



Routledge Studies in Ethics and Moral Theory

THE AUTONOMY OF NORMATIVITY

LOGICAL AND METAPHYSICAL
INTERPRETATIONS OF THE IS-UGHT-GAP

Singa Behrens



The Autonomy of Normativity

While the Is-Ought Gap has recently been a topic of growing interest, most contributions are firmly fixed on logical, often quite technical accounts of the autonomy thesis. This book defends two complementary autonomy theses—a modal and ground-based thesis—that provide a deeper and more comprehensive understanding of the nature of normativity.

The autonomy thesis is often motivated by claims about the nature of the normative domain and its categorical difference from the non-normative domain. This book develops two novel interpretations of the autonomy thesis, one based on the notion of grounding and the other based on a notion of logical-semantic entailment, developed within a framework of truthmaker semantics. Together these accounts capture best the informal idea that we cannot ‘get’ something normative from the non-normative. The proposal is based on an analysis of what it means to say that certain propositional content parts are relevant to the instantiation of entailment and grounding relations. Moreover, the book relates the autonomy debate to other important metaethical debates, and it offers a more explicit account of the theoretical commitments of an autonomist position. Finally, it develops simple and elegant formal equivalents of the proposed autonomy theses which facilitate the evaluation of structurally complex proposed counterexamples, which have impeded a substantive autonomy debate.

The Autonomy of Normativity will appeal to researchers and graduate students working in metaethics, metaphysics, and philosophical logic.

Singa Behrens is a Post-Doctoral Research Assistant and Lecturer at Bielefeld University, Germany. Her research focuses on moral philosophy, metaethics, metaphysics, and normativity. Her published work has appeared in *Analysis*, *Philosophical Studies*, and *Synthese*.

Routledge Studies in Ethics and Moral Theory

The Self, Civic Virtue, and Public Life

Interdisciplinary Perspectives

Edited by Nancy E. Snow

Moral Agency in Eastern and Western Thought

Perspectives on Crafting Character

Edited by Jonathan Jacobs and Heinz-Dieter Meyer

Unfair Emotions

Their Morality and Blameworthiness

Jonas Blatter

The Moral Philosophy of Iris Murdoch

Mark Hopwood

Morality Collapses

Against the Right and the Good

Stephen Kershmar

How We Blame

A Theory of Moral Responsibility

Bana Bashour

What is the Point of Moral Philosophy?

Edited by Roger Crisp, Ingmar Persson, and Julian Savulescu

The Autonomy of Normativity

Logical and Metaphysical Interpretations of the Is-Ought-Gap

Singa Behrens

For more information about this series, please visit: www.routledge.com/Routledge-Studies-in-Ethics-and-Moral-Theory/book-series/SE0423

The Autonomy of Normativity

Logical and Metaphysical Interpretations
of the Is-Ought-Gap

Singa Behrens



First published 2026
by Routledge
605 Third Avenue, New York, NY 10158

and by Routledge
4 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

*Routledge is an imprint of the Taylor & Francis Group, an informa
business*

© 2026 Singa Behrens

The right of Singa Behrens to be identified as author of this work
has been asserted in accordance with sections 77 and 78 of the
Copyright, Designs and Patents Act 1988.

The Open Access version of this book, available at www.taylorfrancis.com, has been made available under a Creative
Commons Attribution-Non Commercial-No Derivatives
(CC-BY-NC-ND) 4.0 International license.

Any third party material in this book is not included in the OA
Creative Commons license, unless indicated otherwise in a credit
line to the material. Please direct any permissions enquiries to the
original rightsholder.

Funded by the Deutsche Forschungsgemeinschaft (German
Research Council) through the Emmy-Noether Grant on
Relevance (KR 4516/2-1), the European Research Council
(ERC Grant 101040439, REASONS FIRST), and the Library
of Bielefeld University.

Trademark notice: Product or corporate names may be trademarks
or registered trademarks, and are used only for identification and
explanation without intent to infringe.

ISBN: 978-1-041-04140-5 (hbk)

ISBN: 978-1-041-04139-9 (pbk)

ISBN: 978-1-003-62701-2 (ebk)

DOI: [10.4324/9781003627012](https://doi.org/10.4324/9781003627012)

Typeset in Sabon
by KnowledgeWorks Global Ltd.

**For anyone who has ever suffered from a gap between
the world as it is and the world as it ought to be.**



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

Contents

<i>Acknowledgments</i>	<i>ix</i>
Introduction	1
PART I	
The Normative and the Non-Normative	11
1 Constraints on Taxonomies	13
2 Existing Taxonomies	26
3 State-Based Taxonomy	51
PART II	
Logical Autonomy Theses	79
4 Inference Barriers and Modal Gaps	81
5 The Relevant Gap Principle	101
6 The Modal Profile of Authoritative Normativity	128
PART III	
Metaphysical Autonomy Theses	159
7 Ground-Based Autonomy	161
8 Relevant Autonomy Thesis	179

viii *Contents*

9	The Moral of Autonomy	197
10	Conclusion	217
	<i>Bibliography</i>	219
	<i>Index</i>	227

Acknowledgments

The strongest person in the world (in 2003), Hugo Girard, once said: *It's not the weight of the stone. It's the reason why you lift it.* Anyway, the weight of *this* stone was significantly reduced and the stone lifting experience immensely enriched by the support and advice of several people to whom I am deeply grateful.

This book evolved from my dissertation, submitted to Universität Hamburg in July 2022. Over the years of working on it, I have received support from many people. First and foremost, I would like to thank Stephan Krämer for his invaluable advice, mindfulness, and steady support. I could not think of a better supervisor than him, who, after several hours of administrative work and another two hours of intensive discussion of this book, and despite being unwell, expressed his joy in our exchange of ideas. I'm truly grateful for the opportunity to discuss with him. I am also very grateful to Benjamin Schnieder for his constant support, and his encouragement ever since my first Bolzano seminar. I have greatly benefited from being part of the Emmy-Noether Project on *Relevance* at Universität Hamburg and the discussions with and comments from my relevant colleagues Stefan Roski and Martin Glazier. Their feedback was always extremely helpful and I have learned a lot from them. Stefan was a constant companion on my way. I am also greatly indebted, for their company and readiness to discuss ideas, to Roman Heil and Benjamin Kieseewetter. The number of astute thoughts Roman shared with me is only exceeded by the number of scoops of ice cream we ate together. To Benjamin I owe a lot, in particular valuable advice and new perspectives on philosophical and other problems. I also owe special thanks to Birgit Recki and Volker Gerhardt for their support, their faith in me, and the insight that the right amount of decisiveness about how things ought to be sometimes bridges the Is-Ought Gap. I want to thank members of the doctoral seminar and the research colloquium for theoretical philosophy in Hamburg for helpful feedback on drafts that found their way into this book. Beside the ones already mentioned I am grateful, for the valuable feedback I received, to

Ralf Bader, Selim Berker, Daniele Bruno, Kit Fine, Johannes Korbmacher, Hannes Leitgeb, Tristram McPherson, Alex Moran, David Plunkett, Olle Risberg, Toni Rønnow-Rasmussen, Gideon Rosen, Stephan Schmid, Thomas Schmidt, Bart Streumer, Pekka Väyrynen, and Peter Verdée.

In addition, I have greatly benefited from discussions on several occasions at which I had the chance to present parts of the material. These occasions include the Gesellschaft für analytische Philosophie PhD-Workshop ‘Mathematical Philosophy’ in 2019 (Ludwig-Maximilians-Universität München), the ‘Relevance, Grounding, and Normativity’ workshop in 2020 (Universität Hamburg), the European Congress for Analytic Philosophy in 2020 (Utrecht University), the Kongress der Deutschen Gesellschaft für Philosophie in 2021 (Friedrich-Alexander-Universität Erlangen-Nürnberg), the ‘Reductionism in Meta-ethics and the Philosophy of Mind’ workshop in 2021 (Oxford University), the ‘Explanation and Logic’ conference in 2022 (UC Louvain), the Kongress der deutschen Gesellschaft für Analytische Philosophie in 2022 (Humboldt-Universität zu Berlin); as well as the colloquia at Humboldt-Universität zu Berlin, Lunds Universitet, and Universität Bielefeld. I would like to thank the participants for their much-appreciated feedback.

Work on this book would not have been possible without the financial support of the *Deutsche Forschungsgemeinschaft* through the Emmy-Noether Grant on *Relevance* (KR 4516/2-1) and the financial support of the European Union (ERC Grant 101040439, REASONS F1RST). I gratefully acknowledge that support.

Parts of [Chapters 3, 7, and 8](#) of this book are based on previously published articles. [Chapters 7 and 8](#) contain material from my paper ‘No Normative Free Lunch: Relevance and the Autonomy of the Normative Domain’, published in *Synthese* 199, 13163-13186 (2021). [Chapter 8](#) contains material from my paper ‘Don’t mind the Gap: How Non-Naturalist Should Explain Normative Facts’, forthcoming in *Philosophical Studies*. [Chapter 3](#) contains material from my paper ‘A Semantics for Moral Error Theory’, published in *Analysis* 84 (2), 221–230 (2024).

Finally, I want to thank my parents, my sister, and my partner. It is almost impossible to imagine how writing this book could have worked without their support and steady patience.

Hamburg,
March 2025

Introduction

‘This is the way it has been done for years, so this is how it *should* be done’. According to what is often called *Autonomy Thesis* or *Hume’s Law*, this inference is invalid for a principled reason: we cannot derive a normative conclusion from non-normative premises alone. That we cannot get an ‘Ought’ from an ‘Is’ is indeed a familiar truism. As stated, it is also demonstrably false. Even by means of the simplest logical rules we can infer an *Ought*-statement from an arbitrary *Is*-statement and generate a counterexample. Here is an example:

P: Cows have tails.

C: Thus, you ought to keep your promises or cows have tails.

That said, the initial claim has significant intuitive appeal.¹ As one of its labels suggests, it traces back to a famous passage in Hume’s *Treatise on Human Nature* (§3.1.1):

[...] as this *ought*, or *ought not*, expresses some new relation or affirmation, ‘tis necessary that it shou’d be observ’d and explain’d; and at the same time that a reason should be given, for what seems altogether inconceivable, how this new relation can be a deduction from others, which are entirely different from it.

The counterexamples that threaten the letter of the autonomy thesis do not seem to undermine this motivation. However, it has turned out to be intriguingly difficult to provide a satisfactory formulation of the thesis that is immune to those counterexamples. In this book, I defend two novel, complementary autonomy theses that capture the appeal of the initial claim without being subject to alleged counterexamples that refute simple interpretations. I argue that together the two autonomy theses—a *modal* thesis and a metaphysical dependence thesis in terms of *ground*—provide a comprehensive account of the autonomy of normativity. One fundamental

2 *The Autonomy of Normativity*

claim of this book is that some counterexamples to the autonomy of normativity are merely *spurious* because transitions that seem to take us from the non-normative to the normative do not hold *in virtue of* normative content. That is, the normative content is *irrelevant* to the transition. Key to my analysis is a notion of irrelevance that makes fine-grained distinctions with respect to the role that parts of the propositional content play in modal and grounding relations. One important upshot of this account is that a ground-based autonomy thesis is superior to a modal autonomy thesis. It better represents what we have in mind when we consider the claim whether the normative is autonomous with respect to the non-normative.

But *why* would we think that the normative is—in a sense that needs to be specified—autonomous with respect to the non-normative? And why should we care? A first thing to note is that normative and non-normative claims seem to be importantly different kinds of claims. On the one hand, *non-normative* claims like the claim that hydrogen is lighter than oxygen, that cows have tails, and that the moon orbits the Earth are claims about how the world *is*. On the other hand, *normative* claims like the claim that parents ought to take care for their children, that flourishing of human life is valuable, and that Sarah may take the last candy are also claims about how things *ought* to be or *may* be (done). Moreover, that Sarah may take the last candy does not settle whether she takes the last candy; likewise, that she takes the last candy does not settle whether she may take it. In light of that contrast, transitions from one kind of claim to the other might seem unwarranted. Consider, for example, the non-normative claim that children are highly dependent on their parents. Suppose from this assumption we draw the conclusion that parents ought to take care for their children. But how can the former guarantee the truth of the latter? According to the autonomy thesis, the non-normative claim can only guarantee the truth of the normative claim if an additional normative assumption is made. Roughly, one might assume that people ought to take care for those who are highly dependent on them. The autonomy thesis entails that without that additional normative assumption, the transition—in this case the inference—is not valid.

The claim that a normative addition is always required, however, faces counterexamples. Before presenting some much-discussed examples, let me address the second question: why should we *care* whether the normative is autonomous? Commonly, normative claims are taken to play a crucial role in our life. What this role is precisely is controversial, but the normative is often described as having a guiding function.² Consider, for example, the debate about climate change. Part of the debate is the question of what ought to be done given scientific insights about rates of sea level rise and the extinction of species. These latter non-normative claims are at least in principle verifiable through scientific, empirical methods.

Normative claims are commonly not taken to be verifiable through the same methods. So, when we aim to argue in favor of or explain normative claims we need to clarify whether normative claims are entailed or fully grounded by the exclusively non-normative. This is not only important for the validity of the relevant arguments or the truth of the relevant explanations, it is also relevant to foundational questions about the nature of normativity.

This book does not provide a comprehensive defense of the claim that the normative *is* autonomous with respect to the non-normative. I am sympathetic to that claim, and I hope to show that this claim has a good chance of being true, but this book takes its starting point from the observation that we have not yet settled what the autonomy of normativity amounts to. My aim is to settle this question, and to show that it is a substantive one, not just a formal one.

As I have indicated in the beginning, a simple interpretation of the autonomy thesis along *logical lines* is demonstrably false. I shall call this interpretation *simple logical autonomy thesis* or SLAT for short:

SLAT: No non-normative premises entail a normative conclusion.

To see why SLAT is false, reconsider the initial disjunctive counterexample. The following argument originates from Arthur Prior (1960)'s seminal paper 'The Autonomy of Ethics'. Take any arbitrary uncontroversially normative and uncontroversially non-normative statement. Prior's example is based on the non-normative statement that tea-drinking is common in England (*D*) and the normative statement that all New Zealanders ought to be shot (*N*).³ It is assumed that the disjunction $D \vee N$ is either normative or non-normative. In any case, we face a counterexample to SLAT. First, supposing that $D \vee N$ is normative, we face a counterexample because the non-normative statement *D* entails the normative statement $D \vee N$ (that is the logical form of the initial example). Second, supposing that $D \vee N$ is non-normative, we face a counterexample because the non-normative statement $D \vee N$ together with the plausibly non-normative statement $\neg D$ entails the normative statement *N*.⁴ Anyway, SLAT is false.

I do not expect the reader to be convinced by examples of that sort that normativity is not autonomous. Indeed, concerns about their significance are precisely what these kinds of counterexamples give rise to. Intuitively, they do not refute what we have in mind when we think about the normative domain as being autonomous with respect to the non-normative domain, but they refute particular explications of that claim.⁵ This observation in turn motivates a distinction between *genuine* counterexamples on the one hand and *spurious* counterexamples on the other hand. Roughly, spurious counterexamples merely refute particular explications

4 *The Autonomy of Normativity*

of the autonomy thesis, while they do not seem to bear on the reasons why we took normativity to be autonomous in the first place. By contrast, genuine counterexamples would show that the normative domain fails to be autonomous. It is important to acknowledge though that at this point, it is an open question whether the distinction between spurious and genuine counterexamples withstands theoretical scrutiny. In case it doesn't, we might need to dismiss another alleged truism.

My aim in this book is to defend a substantial account of autonomy. To that end, I will defend a unified account of spurious counterexamples to autonomy theses in terms of irrelevant content parts. The account allows me to formulate two novel autonomy theses that resist spurious counterexamples. As mentioned above, my account comprises a modal and a ground-based autonomy thesis. The modal autonomy thesis that is in the tradition of the majority of contributions to the autonomy debate⁶ reads as follows: no non-normative premises entail a normative conclusion unless the conclusion's normative content is irrelevant to the entailment. In contrast to entailment, grounding is, roughly, an objective metaphysical *priority* relation among truths. According to the ground-based autonomy thesis that provides an underexplored perspective on the debate, no non-normative propositions fully ground a normative proposition in a normatively relevant way. In order to be able to explain the work that is done by the relevance-constraints, I have to say more about the nature of putative counterexamples like Prior's example.

Basically, there are two strategies to resist putative counterexamples. First, one might object to the classification of the involved statements. That is, one might deny that the statements one began with are non-normative or that the statement one ended up with is normative. Second, one might deny that the example shows anything distinctive about the normative domain and its relation to the non-normative domain. The latter strategy requires a compelling account of what I have called spurious counterexamples. In this book, I will be concerned with both strategies to resist putative counterexamples.

In the first part of this book, *The Normative and The Descriptive*, I introduce and discuss a semantic account of *normative* and *non-normative* content. For ease of presentation, I will refer to the latter also as *descriptive* content. Based on that account, I pursue the first strategy to respond to putative counterexamples and show that some of them are based on false classifications. I argue that a satisfying descriptive-normative distinction requires a fine-grained notion of content that cannot be captured in an *intensional* framework. Intensional accounts of content take necessarily equivalent statements to express the same content. This feature puts the accounts in an uncomfortable dialectical position in the autonomy debate. Part of what is at issue in the debate are entailment relations between

normative and descriptive propositions. On an intensional account, distinguishing between normative and descriptive propositions commits us to the claim that no normative proposition is necessarily equivalent to a descriptive proposition and vice versa. This, however, is a substantial and, as I will argue in the first part of the book, illegitimate assumption when entailment relations between these kinds of propositions are at issue. Instead, I will employ a *hyperintensional* account of content based on truthmaker semantics. The truthmaker framework that is developed in a recent series of publications by Kit Fine (2017a, b, c) allows us to distinguish between necessarily equivalent propositions by distinguishing between different wholly *relevant* and thereby *exact* ways for them to be true. The account's fine-grained tools provide the basis to deal with putative counterexamples where, *prima facie*, descriptive premises entail a normative conclusion but where the classifications do not withstand further scrutiny. The following two examples illustrate misclassifications of the relevant kind:

- (**Truth-Teller**) Everything Alfe says is true. Alfe says that it ought to be the case that everyone is sincere. *Thus*, it ought to be the case that everyone is sincere.⁷
- (**Knowledge**) Sarah knows that she ought to help her friend. *Thus*, Sarah ought to help her friend.

On the account I am going to propose, the seemingly descriptive first premise in each of the examples is classified as normative. In each of the examples, at least one way for the first premise to be true contains normative components, which is why the premise is classified as normative. Thus, the inferences do not provide counterexamples. These considerations, however, do not capture all kinds of putative counterexamples. The second and third part of this book, *Logical Autonomy Theses* and *Metaphysical Autonomy Theses*, are concerned with two pertinent strands along which the autonomy thesis can be interpreted: along logical and along ground-theoretic lines. Both interpretations face spurious counterexamples that cannot be dealt with merely based on a misclassification strategy. I argue that existing accounts to deal with these spurious counterexamples are unsatisfying. The accounts fail to draw compelling distinctions between spurious and genuine counterexamples. As a result, they often turn the autonomy thesis into an overly weak claim. It is a significant feature of my account that my autonomy theses specify conditions under which putative counterexamples would refute the autonomy of normativity. Moreover, my account explains what distinguishes these examples from spurious counterexamples that are compatible with the autonomy of normativity. Let me state two additional spurious and two putative genuine counterexamples (in that order), to provide the reader with a broader picture and to

6 *The Autonomy of Normativity*

sketch essential features of my autonomy theses and the analysis of counterexamples I develop in this book:

- (Conditional Obligation) All undertakers are Church Officers. Thus, if Church officers ought to be reverent, then undertakers ought to be.⁸
- (Non-Existence) Mary does not exist. Thus, it is not the case that Mary ought to go to the meeting.⁹
- (Promise) Smith promised to pay five dollars. Thus, Smith ought to pay five dollars.¹⁰
- (Leibniz-Euthyphro) God wills this act. Thus, this act is good and just.¹¹

On the account I am going to defend, a counterexample refutes the autonomy of normativity if and only if the relation between the descriptive and the normative statements holds *in virtue of* the normative content of the conclusion or grounded statement. By contrast, spurious counterexamples hold merely in virtue of non-normative content of the conclusion or grounded statement. The solution that my approach thereby provides is particularly transparent in the disjunctive case. D entails and grounds $D \vee N$ merely in virtue of $D \vee N$'s non-normative content; its normative content is irrelevant. Moreover, the truthmaker-based approach reveals that all spurious counterexamples are eventually based on the same mechanism even when a similar structure is less transparent like in (Conditional Obligation) or in (Non-Existence). While extant accounts have sometimes made a similar claim—e.g., the truthmaker-based approach in (Fine 2021)—I argue that my account improves the distinction between spurious and genuine counterexamples. It avoids the implausible consequence that examples where the normative *and* the descriptive content is relevant to the transition do not count as genuine counterexamples.¹² Note that if examples like (Promise) and (Leibniz-Euthyphro) were valid, the normative content of the conclusion would be relevant to the transition, which is why (if they were valid!) they would undermine the autonomy of normativity.

Another attractive feature of my account is that both autonomy theses have simple and elegant implementations within a formal framework that is based on the key notion of a state. The modal autonomy thesis is equivalent to the claim that normative states (in a model) do not *exactly exclude*, i.e., are not relevant as a whole to the exclusion of, descriptive states. The ground-based autonomy thesis is equivalent to the claim that no descriptive state generates a normative state (in a model). The formal implementations allow us to diagnose that logical complexity on the propositional

level is the *source* of many spurious counterexamples. Before I state how the chapters to follow proceed, let me briefly make two clarifications.

First, every autonomy thesis concerning the normative domain must invoke a taxonomy classifying some entities (statements, propositions, etc.) as normative and others as non-normative.¹³ As I am using it here, ‘normativity’ is a label for a variety of phenomena that are at issue when we use terms like ‘value’, ‘good’, ‘requirement’, ‘ought’, ‘obligation’, and ‘reason’.¹⁴ These terms play a crucial role in philosophical disciplines such as ethics, epistemology, legal, and political philosophy. However, whether there is a unified phenomenon shared by whatever is picked out by those terms and their cognates is controversial.¹⁵ I side with those who assume there to be a shared phenomenon to which we refer by the term ‘normativity’. My discussion is phrased in terms of normativity thus understood.¹⁶ However, the view I develop is not committed to this claim. If there were no unified domain of normativity, both autonomy theses can be restricted to subdomains. Indeed, my account specifies the conditions that arbitrary subject-matter-based domains must satisfy to count as autonomous in the relevant senses. Turning to the non-normative, one might think that normative statements contrast with natural statements as figuring in the natural sciences. This reading, however, is too narrow. As the (Leibniz-Euthyphro) examples illustrates, I follow a common assumption in the debate in taking supernatural statements to be part of the non-normative domain. The same is true for non-normative statements from scientific areas other than natural sciences. The non-normative is everything that is not part of the normative. As stated, I shall call the non-normative domain *the descriptive domain*.¹⁷

The nature of the outlined distinction is itself controversial. In the first chapter of this book, I will state more carefully what kind of commitment is required in order to take part in the autonomy debate. It is commonly agreed that the distinction can be made with respect to language and concepts. Beyond that there is hardly any consensus. Painting with a fairly broad brush, we might say that the metanormative debate between realists and antirealist is about whether the distinction concerns merely language and concepts or whether it also concerns parts of the world like normative facts or properties.¹⁸ For the time being, it suffices to acknowledge that metanormative views which deny that normative statements play a representative role usually accept some minimal notion of truth-value assignment to account for normative statements in, for example, inferential contexts.¹⁹ On that basis, the autonomy thesis can be evaluated.

Second, this book is concerned with broadly logical and metaphysical autonomy theses. However, one might think of alternative ways in which the autonomy thesis can be understood that do not or merely indirectly

concern me here. One of them is the Moorean one according to which normative concepts cannot be analyzed in descriptive terms.²⁰ Another is an ontological one according to which normative facts and descriptive facts are different in kind.²¹ The last one I shall mention here is an interpretation along epistemic lines.²² Roughly, it says that thinkers are not justified in reasoning from purely descriptive assumptions to normative conclusions. Some of the interpretations plausibly bear on the ones I am concerned with. Yet, to which extent they provide satisfying interpretations of the autonomy thesis will not concern me here. For reasons that will be made explicit in due course, my main interest is in a ground-based autonomy thesis and its relation to and bearing on a more traditional perspective on autonomy in broadly logical terms.

The chapters to follow are divided into three parts. **Part I**, *The Normative and The Non-Normative*, is concerned with a precise distinction between the normative and the descriptive in terms of truthmaker semantics. To that end, **Chapters 1** (*Constraints on Taxonomies*) and **2** (*Existing Taxonomies*) introduce desiderata for a satisfying taxonomy based on a careful discussion of existing accounts. In **Chapter 3** (*State-Based Taxonomy*), I introduce the Finean truthmaker semantics and an analysis of normative content in terms of that account. I apply the truthmaker-based taxonomy to complex proposition that play an important role in the autonomy debate and that are particularly hard to assess. This allows me to solve issues resulting from putative counterexamples based on misclassifications. In **Part II**, *Logical Autonomy Theses*, I develop my preferred modal autonomy thesis. **Chapter 4** (*Inference Barriers and Modal Gaps*) provides the ground by exploring and criticizing existing logical autonomy theses. In **Chapter 5** (*The Relevant Gap Principle*), I develop my interpretation of the modal autonomy thesis according to which no descriptive premises entail a normative conclusion unless the conclusion's normative content is irrelevant to the entailment. I show that this proposal is equivalent to a simple and elegant claim about states within the truthmaker framework. To conclude **part II**, in **Chapter 6** (*The Modal Profile of Authoritative Normativity*), I consider the option that subdomains of the normative domain might fail to meet the autonomy condition. I argue that a restricted autonomy thesis that focuses on authoritative normativity is still an interesting form of normative autonomy that is compatible with this result. In **Part III**, *Metaphysical Autonomy Theses*, I endorse a ground-theoretic construal of the autonomy thesis. In **Chapter 7** (*Ground-Based Autonomy*), I provide arguments in favor of the superiority of a ground-based approach and critically discuss an existing ground-based account. In **Chapter 8** (*The Relevant Autonomy Thesis*), I develop and defend my ground-based proposal according to which no normative proposition is grounded in descriptive propositions in a normatively relevant way. Finally, in **Chapter 9**

(*The Moral of Autonomy*), I tie up the ends. I demonstrate how the two complementary interpretations interact and thereby substantially improve our understanding of autonomy and normativity.

Notes

- 1 The technical use of the term ‘autonomy’ originates from [Rynin \(1957\)](#).
- 2 I will address relevant metanormative questions in due course. For a helpful survey over important controversies, see [van Roojen \(2018\)](#).
- 3 [Prior \(1960, p. 201\)](#). This example is needlessly offensive. However, I decided to present the original in part because accounts that I shall discuss later refer to it and in part because the example is supposed to show that it is valid although the normative disjunct is *false*.
- 4 Another class of examples are necessarily false premises that entail arbitrary normative conclusions, e.g., $D \wedge \neg D \vdash N$, and arbitrary non-normative premises that entail necessarily true conclusions, e.g., $D \vdash N \vee \neg N$.
- 5 Many have argued that the examples are peculiar. For example, [Shorter \(2010\)](#) claims that their conclusions are ‘ethically futile’. More recently, [Fine \(2021, p. 895\)](#) calls such examples ‘cheats’.
- 6 Broadly logical autonomy theses in some form have been endorsed or defended in, e.g., [Prior \(1949\)](#), [Hare \(1952\)](#), [Jackson \(1974\)](#), [Karmo \(1988\)](#), [Pigden \(1989\)](#), [Schurz \(1997\)](#), [Russell and Restall \(2010\)](#), [Singer \(2015\)](#), [Fine \(2021\)](#), and [Russell \(2021\)](#). They have been rejected in some form in, e.g., [Rynin \(1957\)](#), [Foot \(1958\)](#), [Prior \(1960\)](#), [Searle \(1964\)](#), [Nelson \(2007\)](#), and [Maitzen \(1998, 2010\)](#).
- 7 [Karmo \(1988, p. 253\)](#).
- 8 [Prior \(1960, p. 203\)](#). A conditional example that is less artificial originates from [Jackson \(2013\)](#): The descriptive statement that Jane paid for an abortion entails the normative statement that if anyone who has paid for an abortion has done something morally wrong, Jane has done something morally wrong.
- 9 [Fine \(2021, p. 897\)](#). A similar structure is instantiated, if we assume the *Ought-implies-Can Principle*: Sarah cannot go to the meeting. Thus, it is not the case that Sarah ought to go to the meeting.
- 10 This is a curtailed version of the example in [Searle \(1964\)](#).
- 11 See [Leibniz \(1989, p. 59\)](#). Leibniz is concerned with *because*-claims. *Because*-claims concern me in part III of this book.
- 12 Interpretations of the autonomy thesis that are immune to what we intuitively take to be genuine counterexamples are in danger of becoming overly weak claims.
- 13 I carefully assess this apparently trivial requirement in the first chapter.
- 14 See [Finlay \(2010\)](#) for a helpful survey of recent use of the term and ways to understand it.
- 15 Some have argued that we should fundamentally distinguish between, for example, the *evaluative* and the *directive* ([Kirchin 2013](#)) or between *favoring* and *mattering* ([Finlay 2006](#)).
- 16 I will set aside issues that result from construing the normative domain broader than I am conceiving of it here. For example, one might think that notions like ‘truth’ or ‘validity’ are evaluative or normative notions, see, e.g., [Horwich \(2018\)](#) and [Urmson \(1953\)](#). A sentence S entails ‘It is true that S ’. This might pose a threat to autonomy theses concerning the normative domain more broadly

10 *The Autonomy of Normativity*

construed. Notice, however, that the autonomy theses that I am going to defend can be put in general terms. If one wishes to hold on to the autonomy of that domain, one can do so by assuming that states that make the expressed propositions true, respectively, are distinct, while arguing that they meet the conditions that will be outlined in due course.

- 17 One might wonder whether the distinction applies to, for example, inquisitive propositions. Since those propositions do not play a role in the autonomy debate, I shall set them aside. However, there is a natural way to extend the classification that I am going to present to those propositions as well.
- 18 How to divide the metanormative field is carefully assessed in, e.g., [Schroeder \(2009\)](#).
- 19 That the autonomy debate can be generalized to non-cognitivist positions is assumed in, e.g., [Mavrodes \(1968, pp. 58–59\)](#) and [Singer \(2015, p. 194, fn. 4\)](#).
- 20 See [Moore \(1903\)](#).
- 21 See, for example, [Enoch \(2010\)](#).
- 22 For a view that is sympathetic to the epistemic reading, see [Maguire \(2017\)](#).

Part I

The Normative and the Non-Normative



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

1 Constraints on Taxonomies

Preliminaries

A taxonomy, in the sense I will be concerned with, is a classification of objects in the interest of structuring philosophical discussions.¹ So, I need to clarify which objects should be classified. The philosophical discussion that I am interested in in this book is the normative autonomy debate. It is notoriously difficult, however, to even state the idea behind normative autonomy without committing oneself to a claim about the objects that make up the normative and the descriptive domain. So, I help myself with the following schema: we cannot get a normative ϕ from purely descriptive ϕ s. Candidates to instantiate ϕ that have been and might be considered depending on what kind of autonomy thesis one is interested in are sentences, formulas, judgments, propositions, facts, properties, etc. This book is concerned with explications of and relations between *logical* and *meta-physical* autonomy theses. With respect to the target of my inquiry, I focus my attention on taxonomies imposed upon propositions.² At this stage, I stay neutral with respect to the nature of propositions. That includes but is not exhausted by questions of whether propositions have sentence-like structure or how fine-grained propositions are individuated. Propositions arguably provide a common ground for metaphysical and logical interpretations of the autonomy thesis. They allow for a theoretically fruitful comparison between different approaches. This focus contrasts with a mainly epistemic focus, for example, on judgments, or a mainly formal focus, for example, on formulas.³ I will refer to the proposition expressed by a sentence S by enclosing it in square brackets.

It is noteworthy that the focus on *propositions*, in particular *truth-apt propositions*, is already a commitment in normative discourse. Non-cognitivist views deny that normative statements or judgments express genuine normative propositional content.⁴ Accordingly, normative statements as such are not true or false or express a proposition in a proper sense. However, it is important to acknowledge that one need not take the commitment to truth-apt

propositions as metaphysically significant as cognitivists do. A *minimalist* view on propositions and truth makes a taxonomy imposed upon propositions available to metanormative views that take descriptive and normative statements to differ already in their functional role.⁵ According to the minimalist position on propositions, by tacking ‘the proposition that’ onto a sentence one does not make any additional ontological commitments to propositional entities, but generates a useful singular term.⁶ This guarantees the taxonomy’s broad availability. That being said, at some point, minimalism in favor of a shared vocabulary comes to its limits. This is the case in particular for metaphysical autonomy theses. In due course, I will indicate points of departure that strike me as most important.

The next step is to fix the suitable kind of categories. The topic of the inquiry determines at least two of them: the *normative* and the *descriptive* category. Before turning to particular features that propositions need to possess in order to count as normative or descriptive, respectively, one might ask which *kind* of feature it is that makes propositions belong to a particular category. In this book, I shall be concerned with features of propositions defined with reference to a proposition’s topic or *subject matter*. So, I am interested in the contrast between propositions that are—in a sense to be specified—*about* the world’s normative aspects and those that are *about* the world’s non-normative aspects. I shall call these categories *subject-matter*-based categories.

In order to avoid any misunderstanding, let me elaborate on the notion of subject matter as I am using it here. A proposition is about a subject matter if its truth conditions relate to how worlds are with respect to the relevant topic. By contrast, it does not suffice to *mention* a particular subject matter to be in the relevant sense about it. To get a better grasp on the distinction, consider the following examples: since Plato, philosophers engage in a concept that is nowadays called ‘grounding’. The proposition expressed by this sentence is about the history of philosophy. Contrast this with: history of philosophy is a topic that can be studied in a library. The proposition expressed by this sentence mentions but is not about the history of philosophy. While a detailed analysis has to wait until I present my taxonomic account, the truth conditions of the former but not the latter proposition require the history of philosophy to be in a particular way. Similarly, the proposition that normative questions are discussed in school mentions normativity but is not in the relevant sense about normativity. By contrast, the truth conditions of the proposition that there is a moral obligation to help others require the world’s normative aspects to be in a particular way.

To further improve our grasp on *subject-matter*-based categories, consider other kinds of features of propositions that might come to mind in order to categorize them. For example, one might think that structural features of propositions matter to their classification. A structural category

might comprise, for example, all conjunctive or all universally quantified propositions. Accordingly, a proposition would count as normative iff it contains normative components.⁷ In setting aside structural categories, I do not deny that syntactic criteria ultimately succeed. My focus on subject-matter-based categories leaves it open that there is a robust connection between propositions that contain normative components and propositions that are *about* normativity. Yet, this connection cannot be presupposed but requires an argument.⁸ Another kind of feature is an *extrinsic* feature. An extrinsic category might comprise all propositions Ann had conceived of or all propositions expressed by the Old Testament.⁹ One might think that propositions count as normative in virtue of extrinsic features such as being commanded by God. Again, this assumption seems only legitimate if there is a robust connection between the extrinsic feature and how worlds turn out to be normativity-wise.

Let me take stock: a taxonomy that is relevant in the context of the autonomy debate distinguishes between *propositions* that are *about* normative and those that are *about* non-normative matters. It provides the ground for logical and metaphysical autonomy theses that are concerned with a feature inherent in normativity.

Taxonomic Desiderata

In this section, I consider five desiderata for a taxonomy that introduces subject-matter-based categories. I shall call them ‘desiderata’ rather than ‘constraints’ because they intuitively count against a taxonomy if they are not satisfied, yet this fact might be outweighed. In contrast, the violation of the *non-triviality* and the *impartiality* constraint discussed later disqualifies an approach from structuring the philosophical discussion about autonomy. Here is the list of desiderata. I will elaborate on each of them in turn.

- Correctness:** The taxonomy delivers the intuitively correct verdicts with respect to atomic paradigmatic cases.
- \wedge / \vee -Closure:** The taxonomy entails that the normative and descriptive category are closed under conjunction and disjunction, respectively.¹⁰
- Orthodoxy:** Moral laws or normative principles involving descriptive application conditions are normative according to the taxonomy.
- Consistency:** The taxonomy characterizes relevantly similar propositions in the same way.
- Specificity:** The taxonomy is fine-grained enough to account for relevant differences between otherwise similar propositions.

First, consider *Correctness*. This desideratum does not need much justification. Intuitively, the category of many atomic propositions is relatively straightforward. The atomic proposition [Atom *a* has spin-down] is paradigmatically descriptive.¹¹ By contrast, the atomic proposition [Sarah ought to help Ann] is paradigmatically normative. A taxonomy that does not deliver the corresponding verdicts would be revisionary. The fact that a taxonomy classifies paradigmatic cases correctly counts in favor of that taxonomy.

Second, consider \wedge / \vee -*Closure*. If we have two normative propositions, then stating one-*and*-the-other and also stating one-*or*-the-other is to state normative content. The informal motivation is that neither listing normative matters nor counting normative options quits the normative domain. Thus, conjunctions and disjunctions of normative propositions are normative. To consider one example: that Sarah ought to help Ann or she may call for help is plausibly a normative proposition. The descriptive case is analogous. Note that \wedge / \vee -*Closure* is silent about what happens if we conjoin or disjoin a normative *and* a descriptive proposition.

Third, consider *Orthodoxy*. Some taxonomies introduce non-binary distinctions. Accordingly, the normative and the descriptive domain together are not exhaustive. The taxonomies exclude, for example, mixed propositions, i.e., propositions with normative and descriptive components, from both categories. However, taxonomies that generally exclude mixed propositions from both categories cannot account for the normative status of moral laws like ‘If somebody is in distress and help is easily given, then it is obligatory to help’. The reason is that the antecedent of the conditional is descriptive, while the consequent is normative. The taxonomic result that moral laws are not normative, however, cannot do justice to normative inquiry.

Last, consider *Consistency* and *Specificity*. The two desiderata are supposed to guarantee a systematic and coherent account. To see the aim of *Consistency*, consider the following case: the taxonomic status of propositions like [Sarah ought to help Ann or $2+2=4$] and [Tom ought to help Ann or $2+3=4$] is hard to assess. However, intuitively they should belong to the same category in any plausible descriptive-normative distinction. On the other hand, some proposition might be apparently similar, yet differ in relevant aspects. An example that concerns *Specificity* can be generated by the knowledge- and the belief operator. The propositions [Ann believes that Sarah ought to help her] and [Ann knows that Sarah ought to help her] differ at least with respect to their normative implications. The latter but not the former entails that Sarah ought to help Ann. A taxonomy should observe and account for this difference. Note that in the application of the desiderata much hinges on what counts as relevantly (dis-)similar. Yet, this is a feature not a bug. The desiderata can be satisfied

by providing a convincing story about which propositions should and which shouldn't be treated similarly.

To sum up, each of these desiderata might be (partially) rejected or restricted. However, I take it that these maneuvers require at least justification. Satisfying the above desiderata counts in favor of a taxonomy.

Non-Triviality Constraint

While sorting things can be fun, it is a non-trivial task—but the sorting becomes trivial more easily than one might expect. In this section, I focus on results established by Lewis (1988) and Humberstone (1996) concerning the trivialization of taxonomies (*Humberstone-Lewis-results* for short). My aim is to show that although the Humberstone-Lewis-results are often acknowledged in the literature, their significance has usually not been fully recognized. So, in this section, I examine the following formal constraint:

Non-Triviality: The categories of the taxonomy are non-trivial, i.e., they are neither empty nor universal (comprising all propositions).

One might think that this is not only a plausible constraint but also should be taken for granted. Whenever a taxonomy entails that all or no propositions are normative, something went wrong. The Humberstone-Lewis-results, however, show that triviality emerges easily. To see how easy a category becomes trivial, let us first try to come up with plausible taxonomic principles for the normative category. Paradigmatic examples of normative propositions are atomic propositions that attribute normative properties. So, granted that we have a set of atomic normative propositions, it seems plausible to assume that all truth-functional combinations of them are normative as well. Indeed, I took closure under conjunction and disjunction to be a desideratum for a taxonomy. Similarly, it seems plausible to assume that the negation of a normative proposition like [Sarah is not permitted to be lazy] is normative. The intuitive appeal of these principles gets additional support from formal accounts of subject matter. The accounts assume that *subject matter* is closed under truth-functional operators.¹² Formally, this taxonomic principle can be spelled out as follows where \mathcal{N} is the normative category:

Composition: If $S_1, S_2, \dots \in \mathcal{N}$, then $\#(S_1, S_2, \dots) \in \mathcal{N}$, where $\#(S_1, S_2, \dots)$ is any truth-functional compound of S_1, S_2, \dots ¹³

However, this taxonomic principle does not suffice to cover the whole domain of propositions, for propositions are not limited to truth-functional compounds. Consider, for example, the plausibly normative proposition

that all people in public positions ought to be reliable. What seems to make this proposition normative is that it makes a normative claim. It is a non-trivial task to make this intuitive assessment formally precise. A first attempt might be the following: if a normative proposition is true, then it determines part of normative reality. One might think that determination, in turn, can be understood in terms of entailment. Thus, the taxonomic principle would be: whenever a proposition entails a normative proposition, it is normative. And formally:

Converse Entailment: If $S \in \mathcal{N}$ and S' entails S , then $S' \in \mathcal{N}$.¹⁴

On the face of it, both taxonomic principles have some intuitive appeal. It is easy to see, however, that if Converse Entailment is understood in terms of classical entailment, the principles together violate the Non-Triviality Constraint: if \mathcal{N} is supposed to be non-trivial, it cannot be empty. Thus, some proposition N is in \mathcal{N} . Given Composition, $\neg N$ is in \mathcal{N} , which in turn shows that $N \vee \neg N$ is in \mathcal{N} . Since $N \vee \neg N$ is necessarily true, it is vacuously entailed by any arbitrary proposition P . Thus, by Converse Entailment, any arbitrary proposition is in \mathcal{N} . Hence, \mathcal{N} comprises all propositions. Consequently, a taxonomy that assumes both Composition and Converse Entailment violates the Non-Triviality Constraint.¹⁵

One might object that Converse Entailment interpreted in terms of classical entailment does not properly capture the idea of making a normative claim. It is a well-known feature of classical entailment that it allows for vacuous entailments like in the above example. Arbitrary P entails, but is in general completely unrelated to $N \vee \neg N$. This illustrates that it does not make sense to capture the pre-theoretical idea of making a normative claim in terms of classical entailment because otherwise, every proposition makes a normative claim. However, the Humberstone-Lewis-results show that even if one interprets Converse Entailment in terms of a stricter notion of entailment that does not allow for the previous inference, the two principles together turn each category into a trivial category. Going through the proof is illuminating: let P be an arbitrary proposition. I show that P is in \mathcal{N} . Since \mathcal{N} is non-empty, there is N in \mathcal{N} . I assume a notion of entailment that allows for conjunction-elimination, disjunction-introduction, and double-negation introduction. By Converse Entailment, $P \wedge N$ is in \mathcal{N} because $P \wedge N$ entails N . By closure under negation, $\neg(P \wedge N)$ is in \mathcal{N} . Since $\neg P$ entails $\neg(P \wedge N)$, $\neg P$ is in \mathcal{N} by Converse Entailment. Again, by closure under negation, $\neg\neg P$ is in \mathcal{N} . A final application of Converse Entailment shows that P is in \mathcal{N} because P entails $\neg\neg P$. Since P was arbitrary, \mathcal{N} comprises all propositions.¹⁶

One might think that this result is not particularly striking. After all, it only shows that the two taxonomic principles cannot be assumed together.

However, it is not unusual that we learn that two things that we initially found plausible are actually incompatible. However, the result is of *enormous significance* to autonomy debates. The far-reaching consequences for an autonomy debate are often overlooked in the literature. It has been acknowledged that Prior's examples, which are based on those principles, show that one of them must go in order to *avoid counterexamples* to the autonomy of normativity. This assessment, however, misses the point. To illustrate, consider, for example, [Karmo \(1988\)](#), who asks his readers to do the following exercise:¹⁷

Exercise: show that a theory having both features [being closed under negation and converse-entailment] will make ethics non-autonomous, in the sense of admitting sound arguments from non-ethical premises to ethical conclusions.¹⁸

Humberstone shows that this exercise cannot be done. If both principles are assumed, the ethical category is empty or universal. Thus, either there are no ethical conclusions or there are no non-ethical premises.¹⁹ Together, the principles trivialize any autonomy thesis.

Moreover, it is important to acknowledge that Converse Entailment is a one-premise variant of a simple autonomy thesis according to which no non-normative premises entail a normative conclusion. This intimate connection between a taxonomic principle on the one hand and an explication of the autonomy thesis on the other hand seems problematic. It shows that a taxonomy meeting the Non-Triviality Constraint not only has to reject one of the two closure principles but is also in danger of trivializing particular autonomy *debates*.²⁰ To see why, note that a non-trivial taxonomy that assumes *closure under negation* must reject Converse Entailment, if the entailment relation satisfies the elimination and introduction rules I used in the proof. To reject Converse Entailment in turn requires to reject a simple logical autonomy thesis based on the same notion of entailment for merely classificatory reasons. Similarly, a taxonomy that assumes *Converse Entailment* must not only provide a convincing story about why negations of normative propositions are not normative, but also, to avoid begging the question, assume different entailment relations at work in the taxonomic principle and in an explication of the autonomy thesis. Otherwise, the taxonomy is not available for non-autonomists. In the next section, I will turn to the availability of taxonomies in the context of the autonomy debate in detail. For now, the Non-Triviality constraint—far from being a trivial constraint—illustrates that a non-trivial taxonomy must not only reject one of two initially plausible taxonomic principles but also suggests that the very project of a taxonomy already bears on the kind of relation considered to figure in a plausible interpretation of the autonomy thesis.

I conclude with the following observation: while Converse Entailment might appear to be just a taxonomic principle, it is a substantial assumption. It entails substantial consequences concerning the relations between the descriptive and the normative domain. In general, if one assumes taxonomic principles based on a notion of entailment one has two options: either one admits that the defined category is a technical category that is autonomous by design, or one provides an interpretation of autonomy that does not make use of the same notion of entailment as the taxonomic principle does. In the next section, I turn to a methodological constraint on taxonomies that puts pressure on the strategy of defining categories that are autonomous by design.

Impartiality Constraint

In the previous section, it emerged that non-trivial taxonomies are *formally* constrained. A category cannot be both closed under negation and under converse entailment without being empty or universal. I shall now turn to a *methodological* constraint on taxonomies. It has been argued that to serve our methodological purposes, taxonomies must be sufficiently neutral with respect to metanormative and normative theorizing. In particular, it became clear that a category that is autonomous by design is in danger of begging the question in the context of the autonomy debate. An objection frequently made against existing taxonomies is that they fail to meet such a neutrality constraint.²¹ My aim is to clarify the claim made by the neutrality constraint and discuss to what extent it is reasonable to require taxonomies to be metanormatively neutral. I do so by considering two proposals from the literature that explicate neutrality constraints in a slightly different context. Based on these explications, I argue that neutrality is an overly strong requirement on taxonomies. In response, I propose to adopt an *impartiality* instead of a *neutrality* constraint.

To motivate the requirement that a taxonomy must be sufficiently neutral with respect to metanormative views, consider a quote from [Maguire and Woods \(2017\)](#), where the authors criticize model-theoretic taxonomies:

It is *only if* the essential features of the concept of the ethical can be cashed out in a **neutral** fashion in formal terms that the sort of formal picture [...] would be illuminating in the way that many, including ourselves, hope that Hume's dictum is.²²

The idea is that if a taxonomy already takes sides in a debate, it can no longer serve as common ground for the debate in question. While the neutrality constraint is widely held, it is not stated explicitly what neutrality amounts to in this context. To discuss whether a neutrality constraint

is appropriate, a better understanding of the concept and its application to the autonomy debate are required. To that end, let me consider two proposals from the literature that are concerned with the neutrality of the metaethical discourse with respect to ethical discourse. Applied to taxonomies, I examine whether a taxonomy is neutral with respect to autonomy discourse comprising autonomy theses and their negations. The first explication generalizes a proposal in [Fantl \(2006\)](#) that is based on [Dreier \(2002\)](#):

Neutrality_{NC}: A domain of discourse D_1 is neutral with respect to a domain of discourse D_2 iff D_1 does not commit us to the denial of some D_2 -statement.²³

Let me first clarify how to understand the locution ‘ X commits us to Y ’. [Fantl](#) points out that it can be interpreted in different ways that do not always fall together. He considers two interpretations: (i) in terms of entailment, i.e., X entails Y and (ii) in terms of ideally rational belief, i.e., it is ideally rational to believe Y given the willingness to assert X .²⁴ For my purposes, the differences between the interpretations do not matter. To keep things simple, I adopt the entailment interpretation.

Note that **Neutrality_{NC}** has been criticized for being both too strict and too weak to capture a reasonable form of neutrality. For example, [Clipsham \(2015\)](#) points out that according to the definition, the proposition [There are green peas] is not neutral with respect to normative discourse. That there are green peas commits us to the denial of the arguably normative conjunction [It is permitted to eat green things and it is obligatory not to eat peas]. Thus, the definition seems to be too strong.

Let us grant that the proposal can be adjusted accordingly.²⁵ The definition illustrates that a taxonomy cannot be neutral with respect to metanormative discourse wholesale. For one thing, taxonomies might be liberal enough to allow for theoretically light-weight interpretations of truth and propositions to make room for non-cognitivist metanormative positions. Yet, the very project of a taxonomy imposed upon propositions is incompatible with a *radical* denial of the thesis that normative statements express propositional content.²⁶ Moreover, consider a toy metaethical position according to which the term ‘good’ is synonymous with the term ‘most pleasurable to me’ (*Synonymy* for short). A taxonomy that categorizes propositions like [a is good] as normative and propositions like [a is most pleasurable to me] as *non*-normative plausibly entails the denial of *Synonymy*. This illustrates that it is too demanding to require neutrality. One can hardly require that a taxonomy that is supposed to distinguish between descriptive and normative propositions does not make any difference between these propositions. It seems fair to assume that there is a

distinction to be made, for otherwise, debates that are based on the taxonomy would not even get off the ground.

Enoch (2010) criticizes the previous definition for being too weak. He argues that even if D_1 does not on its own commit us to the denial of a D_2 -statement, it might fail to be neutral. The reason is that D_1 might nevertheless be a *difference-maker* to D_2 in the sense that D_1 together with auxiliary premises entails a D_2 -statement where the D_1 -premises are non-redundant. In response, Enoch proposes the following alternative explication:

Neutrality_{ND}: D_1 is neutral with respect to D_2 iff D_1 is a conservative extension of D_2 , i.e., the extended domain $D_1 \& D_2$ preserves good arguments in D_2 and any D_2 -proposition that is supported by a good argument in the extended domain $D_1 \& D_2$ is supported to the same degree in the original domain D_2 .²⁷

The informal idea is that no proposition in D_2 gets *additional support* from the extended domain that it did not already get in D_2 . Let me explain why this explication of neutrality sheds some light on why neutrality constraints are misguided in their application to a distinction between normative and descriptive propositions. Note that the initial characterization of neutrality relies on a notion of difference-making. A domain of discourse D_1 is neutral with respect to another domain D_2 only if it is no difference-maker with respect to D_2 . It is noteworthy that *relevance* is often characterized in terms of difference-making.²⁸ For something to be relevant to another thing, it needs to make some difference to that other thing. However, it is implausible to assume that the way in which the normative and the descriptive domain are separated does not make any difference to metanormative views in general and in particular to autonomy discourse. To illustrate, consider a view called normative nihilism. At first glance, one can formulate the normative nihilist's view as follows: no normative proposition is true. However, this formulation of normative nihilism is trivially false if the normative domain is non-empty and closed under negation, for one of N and $\neg N$ is true.²⁹ This might be taken to show that the normative nihilist's claim needs to be reformulated. However, it illustrates that a taxonomy can make a difference to metanormative views. Given neutrality, the distinction between normative and descriptive propositions would turn out to be irrelevant to the autonomy debate. This strikes me as a quite implausible assumption to begin with.

What seems to be required is a form of *impartiality* of taxonomies with respect to the debate in question, e.g., the autonomy debate and the supervenience debate. A taxonomy might entail particular results in the autonomy debate, for example, that an autonomy thesis is ultimately true (false)

or that particular examples fail to be counterexamples. This, indeed, is a crucial point: each taxonomy will entail which propositions are apt for putative counterexamples. After all, this is the reason why one needs a taxonomy in the first place. Yet, the classification of propositions must not be biased. *Taxonomic principles* that are too close to particular explications of the autonomy thesis, for example, because they exploit the same notion of entailment, violate an impartiality requirement. The taxonomy's classificatory principles must be reasonable assumptions for both sides of the debate. The difference between *neutrality* and *impartiality* in this context is that an impartial taxonomy may eventually entail consequences concerning autonomy and thereby fail to be neutral, but may not *explain* the distinction between the normative and the descriptive in terms of theses at issue. To take this distinction into account, I propose to adopt the following impartiality constraint:

Impartiality Constraint: The taxonomy is impartial with respect to the philosophical debate in question. That is, the *explanatory basis* of the taxonomy does not privilege one side of the philosophical debate due to the fact that *it* contains elements that are unacceptable to one side of the debate.

The impartiality constraint is compatible with the result that the taxonomy, possibly together with auxiliary premises, entails results that are part of the philosophical debate in question. However, the classifications may not be *explained* in a way that presupposes this result. To see the difference, consider a taxonomy that states that all and only atomic, non-negative propositions that involve normative concepts are normative. This taxonomy might entail that the normative domain is closed under converse entailment, but it does not presuppose this principle to define the taxonomic status of propositions. By contrast, a taxonomy that states that all and only propositions that have normative implications are normative presupposes closure under converse entailment to define the taxonomic status of propositions. However, it is incompatible with the aim of structuring philosophical debates to presuppose principles at issue in the debate.

Summary

In this chapter, I have discussed some basic questions concerning the role and form of a taxonomy that distinguishes between the normative and the descriptive domain. I have argued that a satisfying taxonomy must meet several desiderata. It needs to assess paradigmatic cases correctly; in particular, it needs to classify moral laws with descriptive antecedent conditions as normative and satisfy plausible closure conditions. Moreover, the

taxonomy needs to be fine-grained enough to account for relevant differences, while treating relevantly similar propositions similarly. I have also emphasized that taxonomies are formally constrained in order to avoid triviality. A category cannot be both closed under negation and converse entailment. I have argued that this insight bears on how to understand the role of a taxonomy in the autonomy debate. One option is to adopt a less substantial notion of autonomy where autonomy is more like a formal feature of a technical category. This position makes room for stipulations, but might fail to reveal something informative about normativity. Otherwise, one should be happy to reject Converse Entailment. This suggests that a simple logical autonomy thesis is not at the heart of the autonomy debate. I also argued in favor of a methodological constraint that requires an impartial rather than a neutral taxonomy. If my arguments are correct, neutrality is too demanding and threatens to render the taxonomy irrelevant. Impartiality, in contrast, focuses on the taxonomy's explanatory basis. I conclude that taxonomies and their presuppositions, far from being mere preliminaries that can be left implicit, play a crucial role in the autonomy debate. After emphasizing the importance of taxonomic presuppositions and their role in the autonomy debate, I shall now turn to existing taxonomies discussed in the literature.

Notes

- 1 See [Humberstone \(1996\)](#).
- 2 Note that this focus does not presuppose that a taxonomy imposed upon propositions is prior to taxonomies imposed upon, e.g., concepts. At this stage, I stay neutral with respect to any priority ordering.
- 3 As outlined in the introduction, an epistemic autonomy thesis might state that normative judgments are not justified if they are based just on descriptive judgments. A purely formal autonomy thesis might state that a normative formula cannot be derived just from descriptive formulas.
- 4 For an overview over metanormative positions, see [van Roojen \(2018\)](#).
- 5 For a helpful discussion of minimalism in the context of metaphysical in contrast to semantic metaethical doctrines, see [Asay \(2013\)](#).
- 6 See [Asay \(2013, p. 215\)](#).
- 7 This criterion requires that propositions are structured. The taxonomy that I eventually propose need not make this assumption. For the current purpose, one might simply consider the sentences that express the relevant propositions.
- 8 I return to a syntactic account in [Chapter 2](#).
- 9 This contrast is particularly relevant for some of the desiderata introduced in the next section. Take, for example, closure under conjunction. The conjunction of propositions that Ann conceived of is not guaranteed to be a proposition that Ann conceived of.
- 10 More formally, a category \mathcal{C} is closed under conjunction (disjunction) iff for any $S, T \in \mathcal{C}$: $S \wedge T \in \mathcal{C}$ ($S \vee T \in \mathcal{C}$).
- 11 Recall that by enclosing a sentence S in square-brackets, I refer to the proposition expressed by S .

- 12 See, e.g., [Yablo \(2014\)](#) and [Fine \(2017b\)](#).
- 13 I adopt the formalization from [Humberstone \(1996, p. 123\)](#).
- 14 See again [Humberstone \(1996, p. 123\)](#).
- 15 [Humberstone \(1996, p. 132\)](#) shows that *closure under negation* together with Converse Entailment suffices to trivialize a category.
- 16 The proof holds in First Degree Entailment.
- 17 For a more recent misunderstanding, see [Maguire \(2015, p. 192, fn. 5\)](#).
- 18 [Karmo \(1988, p. 257, fn. 7\)](#).
- 19 It is assumed that no ethical proposition is non-ethical.
- 20 In the next chapter, I discuss taxonomic accounts from the literature. It will be made explicit which of the two principles the accounts reject, respectively.
- 21 See, for example, [Maguire and Woods \(2017\)](#), [Fine \(2021, p. 5\)](#), and [Clipsham \(2015, p. 816\)](#).
- 22 [Maguire and Woods \(2017, p. 439\)](#); bold emphasizes added.
- 23 ‘On my conception of moral neutrality, a metaethical position fails to be morally neutral if it commits you to the denial of some moral statement’ ([Fantl 2006, p. 29](#)).
- 24 See [Fantl \(2006, p. 28\)](#).
- 25 [Fantl \(2006, p. 38\)](#) proposes to restrict his definition to the denial of initially plausible statements.
- 26 Early forms of emotivism like in [Ayer \(1936\)](#) deny that moral statements express propositional content at all.
- 27 See [Enoch \(2010, pp. 143–44\)](#).
- 28 Most recognizably in the debate about causation, see, e.g., [Strevens \(2008\)](#). For applications in the grounding debate, see, e.g., [Krämer and Roski \(2017\)](#).
- 29 See [Maguire \(2015, pp. 208–209\)](#) for detailed discussion.

2 Existing Taxonomies

Normative Components

One simple taxonomic account for classifying *sentences* adopts a syntactic criterion.¹ For example, one might classify a sentence as normative if and only if it contains a normative expression. However, one and the same sentence can be used to express normative or descriptive content. The sentence ‘Sarah ought to be here by now’ contains the normative expression ‘ought’ but plausibly has a normative and a non-normative reading. Under the non-normative reading one might want to indicate that Sarah’s train is running late.² To distinguish between both readings, one might think that the component-criterion should be applied to propositions instead. Accordingly, *Normative Components* reads as follows:

NC: A proposition *P* is normative iff it contains a normative sub-propositional component.

Since I am interested in subject-matter-based categories, I have to scrutinize whether NC accommodates all and only propositions that are about normativity. First, note that the taxonomy is incompatible with intensional accounts of propositions, for it requires that propositions are individuated via their sub-propositional parts. Intensional accounts take the sentences ‘This atom has spin-down’ and ‘This atom has spin-down and (lying is morally wrong or it is not the case that lying is morally wrong)’ to express the same proposition, for they are true in the same possible worlds. The component criterion would plausibly distinguish between the two propositions by classifying one as descriptive and the other as normative. This is, of course, not a decisive argument against the taxonomy. Intensional accounts might simply be too coarse-grained to make the relevant distinctions.³ However, normative components neither seem to be necessary nor sufficient for a proposition to have normative subject matter. To see that they are not sufficient, recall that propositions like [Ann believes that Tom

ought to help her] contain normative sub-propositional parts, but do not seem to be normative. They ascribe a belief to a subject. The taxonomy thus fails to provide a criterion to distinguish between structurally similar cases which would be required to satisfy the *Specificity*-desideratum.⁴ Moreover, one might even deny that normative components are necessary. Consider the proposition [What Mary just said is correct] which does not contain normative sub-propositional parts. Yet, suppose that Mary just said that the methods used in Guantanamo are morally wrong. NC suggests that the proposition [What Mary just said is correct] is descriptive, while [Mary just said that the methods used in Guantanamo are morally wrong, which is correct] is normative.⁵ The assessment of such cases is complex, but the mere insistence on normative components cannot do justice to the *Consistency*-desideratum, according to which relevantly similar cases should be treated similarly.

Given its shortcomings, one might wonder why I considered NC in the first place, for it seems overly simplistic. The reason to scrutinize the account is that most existing taxonomies entail relatively coarse-grained distinctions which will become clear in due course. Thus, syntactic criteria provide a useful contrast class. In the end of this chapter, I conclude that a satisfying taxonomy must strike a balance between extremely fine-grained and coarse-grained distinctions so that it avoids the shortcomings of both. The normative component account reveals that normative content is not appropriately characterized if taxonomic status is too closely tied to the surface structure.

Normative Ontological Commitments

According to the next taxonomy, normative propositions incur a normative ontological commitment.⁶ Normative propositions entail the existence of entities that have a normative property. As will become clear, the account yields a fairly narrow notion of normative propositions. I argue that even if this restricted notion is granted, the taxonomy is objectionable for assuming principles at issue in the autonomy debate. Moreover, the taxonomic principle can be generalized so that the normative case is merely one instance among others. As a result, autonomy theses based on this taxonomy do not reveal anything distinctive about normativity.

According to the ontological commitment account, a proposition is normative iff it entails the existence of entities with normative properties.⁷ The proposition that Mary's act was morally wrong entails that something exists that instantiates the property of moral wrongness, namely Mary's act. The application to the descriptive domain seems to work equally well. Descriptive propositions such as that atom *a* has spin-down do not impose any normative ontological commitments. The taxonomy accounts for the

normative nihilist who rejects normativity wholesale because the normative nihilist rejects propositions that would commit her to the instantiation of normative properties. Thus, the normative nihilist can accept the latter but not the former proposition.

Let me now turn to a formally precise version of the account: let C be the class of normative concepts, let $E_w(N)$ be the extension of the concept N in a possible world w . $w \not\models P$ denotes that the proposition P is false in w .

NOC: P is normative iff for all w , if $\bigcup_{N \in C} E_w(N) = \emptyset$, then $w \not\models P$.⁸

NOC says that if the extension of normative concepts is empty in a world, then a normative proposition is false in that world. The taxonomy entails that mixed disjunctions with a normative and a contingent descriptive disjunct are not normative, for disjunctions are true whenever one disjunct is true. The contingent descriptive disjunct in turn can be true in normativity-free worlds. Note that according to the view's own lights, this is a desirable result because the normative nihilist can accept mixed disjunctions in virtue of accepting the descriptive disjunct.

To evaluate the taxonomy, let me first highlight some of its significant features. The taxonomy is closed under converse-entailment. This result straightforwardly follows from the fact that entailment is transitive.⁹ Consequently, the taxonomy is *not* closed under negation.¹⁰ To see why, just note that if P is false in w , then $\neg P$ is true in w . Accordingly, if P is normative, i.e., false in normativity-free worlds, $\neg P$ cannot be normative.

In the literature, the taxonomy has been extensively criticized for conflicting with intuitive classifications and for its limited scope.¹¹ So, I confine myself to listing some of the problematic cases whose descriptive status is particularly striking: mixed material (bi-)conditionals,¹² universal quantifications involving normative concepts, and negations of normative propositions.¹³ It is important to acknowledge, however, that the account does not purport to offer a general account of normative content. Instead, it aims to provide us with a particular interpretation of normative content so that the simple logical autonomy thesis is true with respect to this domain. It focusses on a merely technical normative category to hold on to a simple form of logical autonomy.¹⁴ Due to this strategy, the taxonomy's credibility depends on the theoretical utility of the autonomy thesis it licenses.¹⁵ I will provide two arguments that support the conclusion that this technical sense of autonomy is not informative which in turn renders the taxonomy objectionable.

For the first argument it is important to acknowledge that the commitment criterion in NOC is not bound to *normative* concepts. To see that the criterion is applicable to any arbitrary class of propositions, note that I could have considered the class of *rodent*-concepts instead of normative

concepts. So, I would have been concerned with ROC instead of NOC. I simply have to adjudicate the above pattern. According to ROC, a proposition is a rodent-proposition iff it is false in all rodent-free worlds. This illustrates that the commitment criterion singles out positive content relative to any class of concepts. The corresponding category is closed under converse-entailment and hence entails simple logical autonomy. Thus, rodent-propositions and normative propositions are autonomous in the same sense. Consequently, an autonomy thesis based on NOC does not specify a feature that is particular to the normative or any other domain.

The second argument concerns the general strategy to take normative ontological commitments to be a *unique* feature of normative content. This assumption is at least *prima facie* in tension with supervenience claims. *Strong Supervenience* is the claim that necessarily, if anything x has some normative property F , then there is at least one non-normative base property G such that x has G , and necessarily, everything that has G has F .¹⁶ To see the putative tension, note that it has been argued that Strong Supervenience together with moderate assumptions about closure principles entails that for each normative property there is a necessarily co-extensive (possibly infinite) non-normative base property.¹⁷ The analogous argument can be run for normative truths. This suggests that some (complex) non-normative truths bear a normative ontological commitment.¹⁸ How NOC classifies those propositions depends on whether it is supposed to quantify over metaphysically or conceptually possible worlds. To avoid implausible classifications while allowing for Strong Supervenience it must be assumed that there are conceptually possible normativity-free worlds where the non-normative base propositions are true. Notice that the taxonomy is thereby committed to the denial of conceptual connections between the two domains.

To further strengthen this point, notice that the commitment criterion is quite similar to the Neutrality Constraint scrutinized in the previous chapter. At first glance, it is appealing to say that content that is *not neutral* with respect to normativity is normative. However, it is already a quite strong commitment to assume that a failure of normativity neutrality is unique to normative content. Views according to which normative truths ultimately depend on descriptive truths deny that descriptive content is neutral with respect to normativity. NOC, however, *explains* the classification of propositions in terms of their implications. In other words, the taxonomy presupposes modal autonomy in the explanatory basis and thus fails to be impartial.

Let me summarize the results so far: first, an account that radically restricts the domain to which an autonomy thesis applies must explain in what sense it still characterizes an interesting way in which normativity is autonomous. Second, if the taxonomy is only acceptable to views

that already accept a form of normative autonomy, then something went wrong with respect to the interest of structuring philosophical discussion. The previous section revealed that taxonomies that merely focus on surface structure are not satisfying. In contrast, a taxonomy that makes demands on an ontological level by turning closure under converse-entailment into a taxonomic principle is not impartial and fails to reveal anything distinctive about normativity.

(In-)Variance under Normative Changes

In this section, I examine taxonomies that categorize normative propositions in terms of *variance* under normative changes. I discuss two accounts that assume that a proposition is normative iff its truth-value is sensitive to variations of a normative component in the semantic framework. The underlying thought that *normative structure* is relevant to normative, yet not to descriptive propositions turns out to be promising. However, I argue that a satisfying taxonomy requires a notion of *relevance* that cannot be captured in an intensional framework. First, I discuss the taxonomy in [Russell and Restall \(2010\)](#) that is based on variation of ideal worlds. Second, I turn to the taxonomy in [Singer \(2015\)](#) that is based on variation of norms.

Changing Ideal Worlds

[Russell and Restall \(2010\)](#) characterize the descriptive-normative distinction in terms of *fragility* and *stability* of propositions¹⁹ under variation of normative models. The basic idea is that normative propositions depend on a kind of normative structure in a model in a way that descriptive propositions do not.²⁰ Before introducing the taxonomy, let me flag that Russell and Restall's account provides us with an interesting way to establish *general* inference barriers between two classes of propositions. So, at this stage, my criticism will be limited to the question of whether their categories provide us with an informative way to capture the descriptive-normative distinction.

The taxonomy is based on models for deontic logic and is situated in a quite technical framework. However, the intuitive idea is easy to grasp. The truth-value of propositions involving deontic operators depends on the set of worlds that are *ideal* relative to the actual world. For example, it is obligatory for Mary to help Ann iff in all *ideal* worlds Mary helps Ann. To test whether a proposition is normative one only needs to change the set of ideal worlds and see whether these changes *make a difference* to its truth. If it makes a difference, the proposition is fragile under normative changes, if it doesn't, it is stable. One can then identify

normative propositions with fragile propositions under certain changes and non-normative propositions with stable propositions. The taxonomy is developed in two steps: first, *fragility* and *stability* are defined in an *abstract mode* relative to arbitrary changes. Second, the account defines normatively relevant changes.

Fragility and stability are defined with respect to a set of models \mathcal{M} and a relation R on \mathcal{M} . The taxonomy distinguishes between *stable* propositions where the truth-value is *invariant* under R -changes for every model that satisfies the proposition; and *fragile* propositions where the truth-value *varies* under R -changes for every model that satisfies the proposition.²¹ Note that the two classes of propositions need not be exhaustive, since in general a proposition's truth-value can vary under R -changes with respect to *some but not all* models. Now I turn to the more technical articulation of this idea.²²

R-Stability: A proposition P is R -stable iff whenever a model \mathfrak{M} satisfies P and $\mathfrak{M}R\mathfrak{M}'$, then \mathfrak{M}' satisfies P .

R-Fragility: A proposition P is R -fragile iff whenever a model \mathfrak{M} satisfies P , then there exists a model \mathfrak{M}' such that $\mathfrak{M}R\mathfrak{M}'$, and \mathfrak{M}' fails to satisfy P .

It will become clear how these abstract definitions work in a moment. Before, I need to define normatively relevant changes. The authors consider standard models in deontic logic ($\mathcal{W}, \mathcal{S}, @$): where \mathcal{W} is a set of worlds, \mathcal{S} is a binary deontic accessibility relation, and $@$ is a distinguished actual world. Worlds that are accessible from the actual world ($@Sw$) are morally ideal worlds. Russell and Restall point out that *two* kinds of normative changes are required to capture the whole normative domain. The first relation is *normative extension*. A model \mathfrak{M}' normatively extends a model \mathfrak{M} ($\mathfrak{M}' \supseteq \mathfrak{M}$) just in case you can obtain \mathfrak{M}' from \mathfrak{M} by adding new worlds and extending \mathcal{S} . The second relation is *normative translation*. A model \mathfrak{M}'' is a normative translation of a model \mathfrak{M} ($\mathfrak{M}'' \bowtie \mathfrak{M}$) just in case you can obtain \mathfrak{M}'' from \mathfrak{M} by changing the pairs of worlds related by \mathcal{S} .

The two relations of extension and translation provide us with the notions of extension-fragility (\supseteq -fragility) and translation-fragility (\bowtie -fragility). To get a grasp on these two notions, I will consider some examples. Let O be the ought-operator and P the permission-operator. $O(A)$ is extension-fragile. To see why, note that $O(A)$ is satisfied by a model iff all ideal worlds are A -worlds (see Figure 2.1). Suppose that a model \mathfrak{M} satisfies $O(A)$. Adding new ideal worlds to \mathfrak{M} such that A is false in those worlds, provides us with a normative extension \mathfrak{M}' that does not satisfy $O(A)$ (see Figure 2.2; dotted part). By contrast, $P(A)$ is not extension-fragile.

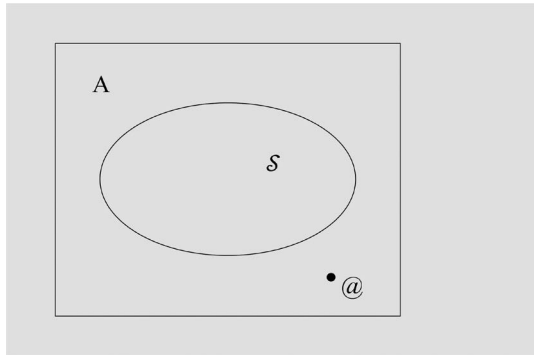


Figure 2.1 Proposition A in a model \mathfrak{M} .

Adding new ideal worlds to a model that satisfies $P(A)$ does not change the fact that there is an ideal world where A is true. This is the reason why extension-fragility does not capture the whole normative domain. However, $P(A)$ is translation-fragile. It is sensitive to the arrangement of \mathcal{S} . Whenever a model satisfies $P(A)$, one can render all A -worlds non-ideal. The translation of the initial model \mathfrak{M}'' does not satisfy $P(A)$ (see Figure 2.2; dashed part). By contrast, $O(A)$ is not translation-fragile. In a model where every world is an A -world no rearrangement of ideal worlds would make $O(A)$ false. Hence, translation-fragility does not suffice to capture the whole normative domain.²³

According to Russel and Restall's taxonomy, a proposition is normative iff it is either translation-fragile or extension-fragile. A proposition is descriptive iff it is translation-stable.²⁴ Propositions containing no deontic




Figure 2.2 A model extension \mathfrak{M}' of \mathfrak{M} (dotted) and a model transformation \mathfrak{M}'' of \mathfrak{M} (dashed).

operators and deontic tautologies turn out to be descriptive according to this approach. It is noteworthy that the taxonomy entails that the normative domain is closed under converse-entailment, yet not under negation.²⁵

One important caveat is that the taxonomy is essentially based on the semantics for deontic operators. However, Russell and Restall argue that their account or at least its mechanism is applicable to the normative domain in general. They argue that we can extend situations in a way that normative but not descriptive propositions are fragile with respect to these extensions. They provide an example that is supposed to show that there is an informal application of the extension-criterion.²⁶ Consider a situation where Alice intentionally hits Bob and nothing else happens (keep this additional condition in mind). The propositions [Alice intentionally hits Bob] and [Alice does something wrong] are plausibly true in this case. Now, we can extend the situation such that Alice and Bob are training for a boxing tournament. According to the argument, that Alice does something wrong is false in the extended situation, while it is still true that Alice intentionally hits Bob. What is more, further extensions such as that Alice knows that Candy threatened to kill all of them, if Alice hits Bob, again change the truth-value of the former proposition.

Let me explain why I think that this informal application is importantly different from the formal account. The formal account determines a *normative* component such that changing it is relevant to normative, yet not descriptive truth-evaluation. In the informal example, however, we *change* the descriptive setting. We change the ‘nothing else happens’-adjunct. This adjunct is relevant to the initial evaluation of the proposition [Alice does something wrong]. We turn the situation from intentional hitting in a *non*-boxing environment into intentional hitting in a boxing environment. The fact that the descriptive proposition is stable under these changes is due to the fact that it is a *merely partial* description of the situation. If we had considered the descriptive proposition [Alice intentionally hits Bob which was unforeseen by Bob] the truth value would have changed. Not all descriptive propositions are stable under all changes. So, what we need is a criterion for variation that affects only normative propositions. At the end of this section, I will raise general doubts concerning the assumption that descriptive propositions are insensitive while normative proposition are always sensitive to variations of *complete* normative structures.

The discussed model-theoretic taxonomy has been criticized from different angles.²⁷ While the details of the objections differ, many of them illustrate a more general point. Before I turn to particular examples of classifications that conflict with an intuitive assessment, let me explain the more general problem. The invariance criterion presupposes that descriptive propositions are stable under variation of normative structure. This is to assume that only normative propositions exclude normative

possibilities. Since locutions like *exclusion*, *stability*, and *fragility* are spelled out in modal terms, these assumptions seem illegitimate in a context where the modal relations between two classes of propositions are at issue. The taxonomy turns a controversial assumption about the normative domain into a taxonomic principle and thus just is the autonomy thesis in a taxonomic guise.²⁸

The issues that result from this general point can be illustrated by particular cases. It has been pointed out that if the set of models is constrained by alethic-deontic principles the taxonomy entails implausible results. For example, if all models satisfy $A \rightarrow O(A)$ and A for some descriptive A , then $O(A)$ is not normative.²⁹ Similar considerations apply to conceptual truths concerning intrinsic normative properties. Maguire and Woods (2017) argue that if it is a conceptual truth that, for example, pain is *intrinsically* disvaluable, then the proposition that Ann's pain is disvaluable does not count as normative according to the taxonomy.³⁰ The normative models cannot be changed in ways that render this proposition false. Note that from the model-theoretic point of view, it is legitimate to assume unconstrained models. However, a semantic account that generally abandons conceptual constraints renders the taxonomy unavailable for positions that assume such conceptual truths. Instead, it provides us with a technical category that is autonomous with respect to another technical category. Whether the categories represent normative and descriptive propositions, respectively, is at least as controversial as the autonomy thesis.

In a last step, I will briefly show that the categories are not applicable to cases that are arguably normative. It is easy to see that necessary truths are not normative according to the taxonomy, for no model falsifies a necessary truth. Consequently, there are no normative necessary truths. Moreover, Vranas (2010) points out that intuitively normative quantifications like [All citizens ought to vote] or [No student may cheat] are neither normative nor descriptive according to Russell and Restall's taxonomy.³¹ The same holds for mixed disjunctions and thereby for material conditionals like [If there is no lift, then Sarah ought to help her grandmother]. Models that do not satisfy the descriptive antecedent thereby satisfy the corresponding disjunction while no changes to the normative structure can render the descriptive disjunct and thus the disjunction false. The strategy to exclude mixed propositions from a descriptive-normative distinction will concern me again, when I turn to autonomy theses. I shall call this general strategy the *exclusion strategy*.

Let me recapitulate. The taxonomy assumes that all and only normative propositions are variant under changes to normative structure. However, it turns out that not all intuitively normative propositions vary under those changes. Moreover, to limit the class of propositions that are intuitively misclassified, the taxonomy must abandon conceptual connections

between the domains. This strategy presupposes normative autonomy and thus renders the taxonomy's availability limited to those who accept the autonomy of the normative domain. After considering a second variance-approach, I argue that this effect is due to the fact that the variance test is based on *a modal notion of relevance*. Changes imposed upon consistent and complete items such as models might turn out to be modally relevant for the descriptive domain. Assuming the contrary is to presuppose autonomy.

Changing Normative Standards

Another variance-based approach is developed in [Singer \(2015\)](#). Singer's taxonomy is based on variation of *norms* representing normative aspects of possible worlds. A proposition's classification depends on whether norm-variation makes a difference to the proposition's truth-value.³² According to the proposal, propositions that are invariant under norm-variation are descriptive; propositions that change their truth-value under norm-variation are normative. The intuitive idea is easy to grasp: which norms hold in a world makes no difference to whether atom *a* has spin-down, but it is relevant to Sarah's praiseworthiness.

To make this idea formally precise, Singer introduces a double-indexed semantic framework. On Singer's account, propositions are modeled as sets of world-norm pair $\langle w, n \rangle$ for $w \in \mathcal{W}$, $n \in \mathcal{N}$ where \mathcal{W} is a set of worlds and \mathcal{N} is a set of norms. Singer describes norms as complete plans for action in any hypothetically possible scenario. Roughly, one may think of the first component w as fixing the descriptive aspects and of the second component n as fixing the normative aspects of possible worlds in the sense of standard possible-worlds semantics.³³ Similar to standard possible-worlds semantics, a proposition P is modeled as a set of world-norm pairs that are compatible with P (denoted by $\langle w, n \rangle \models P$). A proposition is norm-invariant if and only if, for any world, whenever there is some choice of norms such that the world-norm pair makes the proposition true, the proposition is true in that world under any choice of norms. Formally:

Norm Invariance: P is norm-invariant iff $\forall w (\exists n \langle w, n \rangle \models P \rightarrow \forall n \langle w, n \rangle \models P)$.³⁴

If a proposition is norm-invariant, it is descriptive. By contrast, a proposition is normative, whenever there is at least one world such that under one choice of norms the world-norm pair makes the proposition true, while under another choice of norms the world-norm pair makes the proposition false. Note that the categories of norm-variant and norm-invariant propositions are exclusive and exhaustive with respect to propositions that are

made true by at least one world-norm pair. Propositions that can be true either change their truth-value under norm-variation with respect to *some* world-norm pair or they are stable under all variations for *all* world-norm pairs, but not both. As stated norm-(in)-variance is merely a sufficient condition for a proposition's being normative (descriptive). However, for the rest of the discussion, I will call the norm-variant domain 'normative domain' and the norm-invariant domain 'descriptive domain'.

The taxonomy entails that the normative domain is closed under negation.³⁵ Note that in contrast to previous taxonomies, this taxonomy is *not* closed under converse entailment.³⁶ On an account of propositions as sets of worlds, entailment is defined in terms of the subset relation. A proposition P entails a proposition Q iff $P \subseteq Q$. To see that the normative domain is not closed under converse-entailment, consider a mixed disjunction with a contingent normative and a contingent descriptive disjunct. The disjunction is norm-variant because in $\neg D$ -worlds $D \vee N$'s truth-value depends on the truth-value of the normative disjunct N , which is norm-variant by assumption. Thus, descriptive D entails normative $D \vee N$. A noteworthy consequence is that Singer's *taxonomy* is compatible with views according to which descriptive propositions entail normative propositions. As I will show in the next paragraph, the fact that the normative category is neither closed under entailment nor under converse-entailment has been overlooked in some arguments against Singer's view. I will briefly point out where these arguments go astray to obtain an accurate picture of advantages and disadvantages of Singer's taxonomy. However, eventually I argue that Singer's taxonomy is committed to a modal assumption that is controversial and particularly problematic in the context of an autonomy debate.

Maguire and Woods (2017) criticize Singer's taxonomy for counting paradigmatically descriptive propositions as normative and vice versa. I argue that their arguments to establish these results are uncharitable to Singer's *taxonomic* account. I explicitly discuss Maguire and Wood's example that is supposed to illustrate that a paradigmatically normative proposition is classified as descriptive, but the same reasoning applies to the reverse case.³⁷ The paradigmatically normative proposition considered by the authors is [It is wrong to torture Bob]. The authors argue that this proposition is classified as descriptive according to Singer's approach. Their argument goes as follows: according to Singer's taxonomy, a norm that holds at every world-norm pair is norm-invariant and thus descriptive. Note that this is an implausible taxonomic consequence of Singer's account in itself, but for now I shall grant it. For their argument, Maguire and Woods assume that the norm that it is wrong to torture someone just because they have inconvenienced you holds at every world-norm pair. Thus, Singer is committed to classifying this proposition as descriptive. Moreover, the proposition [The only interaction with Bob so far was

one of inconvenience] is assumed to be descriptive. The authors take the two former propositions to entail that it is wrong to torture Bob.³⁸ They conclude that on Singer's account, the conclusion [It is wrong to torture Bob] is descriptive because it is entailed by descriptive premises.³⁹ This reasoning, however, is based on the assumption that the descriptive category is closed under entailment—which it is not. The closure principle is explicitly—and correctly so—rejected by Singer.⁴⁰

To the contrary, Singer's taxonomy classifies the proposition [It is wrong to torture Bob] as normative. It is thus in line with Maguire and Woods's intuitive assessment. The example actually illustrates that Singer's *taxonomy* is compatible with inferences from descriptive premises to normative conclusions. To formulate an *autonomy thesis*, Singer introduces another category that tells us under which circumstances a conclusion counts as *relevantly ethical* relative to a particular argument. A proposition is relevantly ethical relative to a particular argument iff it is norm-variant with respect to world-norm pairs that are not already excluded by the premises. Whether a conclusion is relevantly ethical in an argument depends on its premises. Later, I discuss Singer's autonomy thesis in detail and show that conclusions that are entailed by descriptive premises cannot be relevantly ethical in the context of the corresponding argument. So, one may object to Singer's *explication of an autonomy thesis* on the basis of Maguire and Woods's example, for the conclusion [It is wrong to torture Bob] arguably is relevantly ethical in the context of the outlined argument. Note that this is an objection to the autonomy thesis, but not to Singer's taxonomy. Thus, the examples do not show that paradigmatic normative propositions are classified as descriptive and vice versa.

In the rest of this section, I focus on Singer's taxonomy. It is important to acknowledge that Singer's variance criterion requires that the normative and the descriptive aspects of worlds admit free recombination. Otherwise, if a world's descriptive aspects already fix its normative aspects, one cannot vary the norm-component. Accordingly, intuitively normative propositions turn out to be norm-invariant and thus descriptive. Singer is aware of this fact. In a footnote, he admits that invariantist semantics either need to allow for impossible worlds that differ in which norms are part of the world or allow that normative aspects of a world *are not fixed* by the non-normative aspects.⁴¹

The assumption of free recombination is, however, controversial and particularly problematic in the context of an autonomy debate. [Humberstone \(1996\)](#) convincingly argues that free recombination conflicts with theses like global supervenience according to which worlds that are alike in all descriptive aspects are alike in all normative aspects. Humberstone concludes that a taxonomy that is committed to denying important meta-ethical principles such as global supervenience should be rejected.

Let me briefly sketch Humberstone's argument. In [Humberstone \(1996\)](#), the author develops a taxonomy based on Lewisian subject-matters⁴² that is committed to the assumption that the normative and the descriptive subject matter are orthogonal. A Lewisian subject matter is a partition on the set of possible worlds. Worlds that belong to the same partition cell are alike with respect to a subject matter M . One might think of an equivalence relation \equiv_M between possible worlds that are alike with respect to M . Note that the two components of world-norm pairs on Singer's account can be seen to represent descriptive and normative equivalence classes, respectively. Two subject matters M_1, M_2 are orthogonal iff for any possible worlds w_1, w_2 there is a third possible world w_3 which is equivalent to w_1 with respect to M_1 ($w_1 \equiv_{M_1} w_3$), and equivalent to w_2 with respect to M_2 ($w_2 \equiv_{M_2} w_3$). If the normative and the descriptive subject matter are orthogonal, then for every two worlds, there is a third world that is like the first world in all descriptive aspects and like the second world in all normative aspects. In Singer's framework, this is equivalent to the assumption that for every two world-norm pairs $\langle w, n \rangle$ and $\langle w', n' \rangle$ there is a third world-norm pair $\langle w, n' \rangle$. This gives us free recombination. However, if the two subject matters are orthogonal, two worlds like w_1 and w_3 can be alike in all descriptive, but not in all normative aspects which violates global supervenience. Humberstone dismisses his own taxonomic proposal for being committed to orthogonality.⁴³

I will argue that the conflict emphasized by Humberstone illustrates a more general problem for world-based taxonomic principles. To see why, note that taxonomies such as Singer's taxonomy are based on the assumption that changes to normative aspects of worlds make a difference to the truth of normative propositions, yet not to the truth of descriptive propositions. One may say that according to this assumption normative aspects of worlds are relevant to the truth-evaluation of normative propositions, but not to the truth-evaluation of descriptive propositions. The claim's plausibility depends on the fact that one can separate the world's normative aspects from their non-normative aspects. This leaves the question open which combinations of normative and descriptive aspects are possible. From a purely model-theoretic perspective, one might argue that any logically consistent combination must be considered. However, taxonomies based on unrestricted combination merely allow for quite weak forms of autonomy. To see why, note that on this proposal bachelor-aspects of worlds are considered to be independent from male-and-unmarried-aspects of worlds. As a result, male-and-unmarried-propositions would be invariant under the variation of bachelor-components in our semantics. Accordingly, bachelor-propositions turn out to be autonomous with respect to male-and-unmarried-propositions, which is a very weak sense of autonomy. By setting aside all conceptual and metaphysical constraints,

the model theoretic-perspective might not be able to do justice to a plausible notion of normativity.⁴⁴ Free recombination understood as a conceptual or even metaphysical assumption, however, is controversial, in particular, in light of a debate where the modal autonomy of a domain is at issue.⁴⁵

Another option suggested by Singer is to assume impossible worlds in our semantics. This maneuver might help to mitigate against worries concerning normative supervenience. However, the maneuver does not seem to be compatible with Singer's overall account. To see why, note that usually when it comes to impossible world semantics, entailment is restricted to normal, i.e., possible worlds.⁴⁶ However, if one restricts the entailment relation accordingly, then the taxonomy is incompatible with Singer's autonomy thesis.⁴⁷ Normative conclusions would be norm-variant even with respect to those world-norm pairs that are not already excluded by the premises.⁴⁸ Yet, according to Singer, his autonomy thesis is a *theorem* of any plausible semantics for normative language. I conclude that a taxonomy based on truth-value variation under particular changes that is carried out within a possible worlds framework either yields too coarse-grained distinctions or is committed to controversial modal assumptions such as free recombination.

The Lesson: Normative Difference-Makers

I am now in the position to draw general conclusions from variance-based taxonomies. It is a plausible assumption that the way in which normative aspects of worlds matter to the truth-conditions of normative propositions is quite different from the way in which they matter to the truth-conditions of descriptive propositions. However, the discussion of the previous world-based taxonomies revealed that characterizing the normative domain in terms of variance and the descriptive domain in terms of invariance under normative changes commits us to the assumption that the world's normative aspects are modally independent from the world's descriptive aspects. The reason is that the informal criterion of being relevant to the truth of a proposition is formally spelled out in modal terms. It is assumed that descriptive truths are modally independent from the normative components in a model or a world-norm pair. On the other hand, it is assumed that whenever something is modally dependent on the normative component in a model or a world-norm pair it is normative. These commitments are controversial and are in danger of presupposing exactly those principles that are at issue in the autonomy debate. As a consequence, taxonomies that purport to assume generally acceptable taxonomic principles entail unintuitive results given assumptions about the normative domain that conflict with the taxonomic assumptions. An advantage of the variance-based

taxonomies is that they can prove the autonomy of the technical category they specify. Autonomy is an immediate consequence of a modal variance-criterion. If there is an x such that x makes a difference to Q 's truth-value but not to P 's truth-value, then P cannot entail Q . This advantage, however, comes at a high price. The motivating idea that normative aspects of worlds are relevant to normative propositions in a way in which they are not relevant to descriptive propositions is nevertheless promising. Yet, what seems to be required is a more fine-grained notion of relevance. A hyperintensional notion of relevance will allow me to distinguish between the way in which normative aspects of worlds are relevant to normative propositions and the way in which they are relevant to descriptive propositions even if they modally bear on both. An account that is based on a stricter notion of relevance thus can stay neutral with respect to the question of whether normative aspects of worlds are modally independent from the descriptive aspects, while still assuming a distinction between normative and descriptive propositions based on normative components in the semantic framework.

Taxonomic Relativism

So far, I considered attempts to define normative and descriptive propositions *simpliciter*. However, given that non-paradigmatic cases like mixed disjunctions seem to have features of descriptive and normative propositions, one might think that they should be classified relative to the relevant context. Consequently, a proposition's taxonomic status would be a relation that this proposition bears to some x . One can distinguish relativistic accounts by the ways they fill in the second relatum x . In what follows, I will present three accounts and raise doubts concerning the general strategy. I argue in favor of taxonomic *essentialism*, i.e., the view that the category to which a proposition belongs is essential to it.⁴⁹

Let me first consider three ways to fill in x : the taxonomies assume that propositions or sentences fall into a category relative to an argument, a world, or a speaker.

- Argument-Relativism:** Taxonomic status is a relation that propositions or sentences bear to *arguments*.
- World-Relativism:** Taxonomic status is a relation that propositions or sentences bear to *worlds*.
- Speaker-Relativism:** Taxonomic status is a relation that propositions or sentences bear to *speakers*.

To illustrate the taxonomic consequences, consider the mixed disjunctive sentence 'It is sunny or murder is morally wrong'. According to [Pigden's](#)

(1989) argument-relativism, ‘It is sunny or murder is morally wrong’ is normative relative to the argument \llcorner ‘Murder is morally wrong’ \lrcorner ‘It is sunny or murder is morally wrong’ \gg , but fails to be normative relative to \llcorner ‘It is sunny’ \lrcorner ‘It is sunny or murder is morally wrong’ \gg .⁵⁰ According to Karmo’s (1988) world-relativism, ‘It is sunny or murder is morally wrong’ is normative in a world where murder is morally wrong on a rainy day, but fails to be normative in a world where murder is morally wrong on a sunny day. According to Dreier’s (2002) speaker-relativism, ‘It is sunny or murder is morally wrong’ is normative if uttered by Ann who thinks that it is not sunny. It is not normative if it is uttered by Sarah who thinks that it is sunny. More needs to be said about the details of the accounts to do justice to the proposals. However, I will focus on the general strategy and hence, the details may be set aside.

My aim is to show that relativistic taxonomies answer questions that can be interesting, yet do not help to clarify the autonomy debate. Roughly, they ask whether an argument *depends* on a conclusion’s normative aspects, whether a speaker’s assertion *depends* on normative beliefs, or whether the truth of a sentence actually *depends* on the world’s normative aspects. These questions are worth asking, yet they do not help us answer the question I began with: what is a normative proposition and can we get a normative proposition from descriptive propositions? I will argue that taxonomic relativism is particularly implausible when it comes to the classification of relevantly similar cases. As will become clear soon, the discussion to follow can be understood as a debate about which cases are relevantly similar.

First, consider argument-relativism. Brown (2014) argues that argument-relativism conflates two questions.⁵¹ One question is whether a proposition expressed by a sentence in a given argument is normative, another question is whether the argument’s validity depends on the proposition’s normative status. It is clear that the arguments given above differ with respect to the latter question. \llcorner ‘It is sunny’ \lrcorner ‘It is sunny or murder is morally wrong’ \gg is valid whether the second disjunct of the conclusion is normative or not. By contrast, if one replaces the premise by ‘Murder is morally wrong’ the validity depends on the fact that one disjunct is identical with the normative premise. Brown argues that the fact that something does not depend on *A*’s being *F* does not show that it is not the case that *A* is *F*.⁵² This is a serious objection to argument-relativism. It is a particularly implausible consequence of the approach that if one adds ‘It is sunny’ as a further premise to the argument, then the conclusion ‘It is sunny or murder is morally wrong’ is not normative relative to the modified argument. Hence, the disjunction’s classification differs with respect to the one-premise and the two-premises argument. This seems to violate the *Consistency-desideratum*.

One might think that the same objection applies to world- and speaker-relativism. Yet, world- or speaker-relativists might respond that if the actual truth or assertibility of a sentence *S* does not depend on normative aspects, then this shows that *S* is not normative. This response is closely related to what motivates the normative commitment account. If *S*'s truth depends on a world's normative aspects, then its truth requires the world to be in a particular normative way. Similarly, if a speaker asserts *S* based on normative beliefs, then she commits herself to a normative view. This response, I think, is not satisfying. To the contrary, taxonomies that focus on contingent aspects, i.e., actual ways for a proposition to be true or the actual reasons for asserting it, prevent us from giving an accurate analysis of normative content. To see why, consider the following weather-wrongness-variations:

- (**Rain-Murder**) It is rainy or murder is morally wrong.
- (**Sun-Murder**) It is sunny or murder is morally wrong.
- (**Sun-Donation**) It is sunny or donation is morally wrong.

Let us first turn to world-relativism. Consider a world where it is rainy and murder is morally wrong, but donation is not. This entails the following classifications: (Rain-Murder) is non-normative, (Sun-Murder) is normative in this world. (Sun-Donation) is not normative, if the taxonomy focuses on the actual truth-value in a world.⁵³ It is normative, if the taxonomy focuses on variance of truth-value in a world with respect to normative changes.⁵⁴ An odd consequence of the taxonomy is that the propositions' classification depends on, for example, the actual weather.

Turning to speaker-relativism, a first thing to note is that mixed disjunctions like the weather-wrongness-variations are descriptive if the speaker considers both disjuncts to be true. For example, if a speaker asserts (Rain-Murder) while she believes that it is rainy and that murder is morally wrong she makes a descriptive assertion. This seems to be an implausible consequence. Moreover, in general, it is worth pointing out that the account is a radical response to issues that arise in the context of the taxonomic project. Initially, I tried to categorize propositions as either normative or descriptive. The account suggests that this distinction can only be made when speakers are taken into account. Accordingly, before a proposition can be classified, one must know who expressed or thought about this proposition and based on which assumptions she does. It is a surprising consequence of the account that a proposition's subject matter depends on the corresponding speaker and her background beliefs. This is, of course, no decisive argument against such an account. However, if alternative accounts can do without such a radical move while still providing a coherent and informative distinction, they have a clear

advantage over speaker-relativism. In the next chapter, I propose such a taxonomic account.

As outlined above, some of the former arguments depend on the claim that the weather-wrongness-variations are relevantly similar. The rationale for this claim is that all of them seem to be partially about the weather and partially about moral wrongness. They have something relevant in common, namely their subject matter. Subject matter, however, seems to be largely independent of actual truth.⁵⁵ Consider, for example, the sentence ‘ $\sqrt[3]{19683} = 27$ ’. One can be ignorant about the truth-value and at the same time know that this sentence is about mathematics. This is in line with a principle that Brown endorses:

Super-Mean: The taxonomic status of a sentence supervenes on its meaning.⁵⁶

If this principle is correct, the weather-wrongness-variations seem to belong to the same categories on a plausible descriptive-normative distinction. Whatever the weather may be, the meaning of the sentences remains constant and so does the taxonomic status of the corresponding proposition. I conclude that taxonomic status is essential to propositions. I shall now turn to the question of whether an intensional framework is suitable to capture subject-matter-based distinctions.

Intensional Taxonomies

The autonomy thesis is motivated by the observation that normative propositions are *about* other things than descriptive propositions. As explained in [Chapter 1](#), I take this observation seriously by focusing on subject-matter-based categories. Classically, subject-matter-based taxonomies are developed within an intensional framework based on possible worlds and partitions on possible worlds.⁵⁷ In this section, I argue that intensional taxonomies face serious difficulties due to their coarse-grainedness. While large parts of my criticism are more general, my arguments are developed in detail for [Oddie’s \(2018\)](#) account.

[Oddie’s \(2018\)](#) taxonomy is based on the following definitions: a subject matter *SM*, e.g., number of planets, is a partition of logical space, i.e., the set of all possible worlds, into cells C_1, C_2, \dots , such that for each cell C_i the worlds in C_i are pairwise indistinguishable with respect to *SM*. Each C_i provides a complete possible answer to the question corresponding to the subject matter, e.g., *how many planets are there?*. **NORMATIVITY** partitions logical space in such a way that in each cell N_i all worlds share their normative structure; **DESCRIPTIVITY** partitions logical space in such a way that in each cell D_i all worlds share their descriptive structure. So, the *SM*-content

of an arbitrary proposition P (P^{SM}) is defined as the union of all SM -cells that are compatible with P . The normative and the descriptive content of a proposition P is P 's normativity-content, P^{NOR} , and P 's descriptivity-content, P^{DES} . Note that logically equivalent propositions have the same normative and descriptive content, for they are compatible with the same sets of worlds. This framework provides the basis to define four categories.⁵⁸ First, consider normative and descriptive propositions:

Normative: A proposition P is normative iff $P = P^{NOR}$.

Descriptive: A proposition P is descriptive iff $P = P^{DES}$.

In the former case, the proposition is fully about normativity and in the latter, the proposition is fully about descriptivity. Some propositions have trivial normative (descriptive) content. For example, necessary truths and falsities are compatible with all or no normative (descriptive) cells, respectively. Thus, according to Oddie's account, both necessary truths and falsities are normative and descriptive. In contrast, there are propositions like mixed conjunctions or disjunctions that are neither fully about normativity nor fully about descriptivity. Thus, Oddie defines two further categories:

Fusion: A proposition P is a fusion iff $P \neq P^{NOR}$, $P \neq P^{DES}$, and $P = (P^{NOR} \cap P^{DES})$.

Hybrid: A proposition P is a hybrid iff $P \neq (P^{NOR} \cap P^{DES})$.

Normative-descriptive fusions decompose into their normative and descriptive contents. By contrast, Oddie describes normative-descriptive hybrids as having only trivial normative (descriptive) content, since they are compatible with all normative (all descriptive) cells. For example, a mixed disjunction $D \vee N$ neither excludes descriptive nor normative possibilities and thus is compatible with all normative and all descriptive cells.

Let me now turn to a critical assessment. I discuss three problematic consequences of the account. I start with a consequence that one might be willing to swallow and end with what I take to be the weightiest objection. I shall argue that the issues emerge from the intensional set-up and thus generalize to other intensional accounts.

First, it already became clear that necessary truths are both descriptive and normative. Accordingly, the proposition [Atom a has spin-up or it is not the case that atom a has spin-up] has (trivial) normative content. The fact that its normative content is merely trivial might mitigate against this unintuitive result. However, the distinction between non-trivial and trivial normative content seems more appropriate when considering the

propositions [Murder is morally wrong] and [Murder is morally wrong or it is not the case that it is wrong]. By contrast, the proposition [Atom a has spin-down or not] does not seem to have any normative content at all. Thus, the view entails an implausible account of normative content. One might respond by modifying the account so that necessary truths are neither normative nor descriptive rather than both. Yet, this response would require giving up on the \wedge/\vee -Closure-desideratum.⁵⁹ Moreover, to state that [Atom a has spin-down or not] and [Murder is morally wrong or not] are about the same subject matter requires a departure of any intuitive sense of aboutness. This is no decisive argument, but a cost of the view.

Second, consider a view according to which particular kinds of speech-acts, e.g., promises, generate obligations. For example, suppose that worlds in which Ann promised to ϕ are worlds in which Ann ought to ϕ . In general, if a proposition D entails a proposition N , then D and $D \wedge N$ are necessarily equivalent. On an intensional account, necessarily equivalent propositions share their normative and descriptive content. Thus, if D and N are contingent, then D has non-trivial normative content. However, [Ann promised to ϕ] does not seem to have normative content—at least this is assumed by those views that deny autonomy based on such examples. One might object that the example requires conceptual constraints concerning non-logical terms such as *promise* and *obligation*, while logical space is not constrained in this way. However, one advantage of content-based accounts, in contrast to merely syntactic accounts, is that propositions are classified as normative in virtue of their particular truth-conditions. Setting aside conceptual constraints on normative concepts in the taxonomy project, however, ignores important aspects of truth-conditions of normative content.

Third, I turn to a consequence of the taxonomy that I take to be a decisive reason to reject it. The view entails an unsatisfying account of partial content, particularly of normative partial content. To see why, recall that on the present account mixed disjunctions like $D \vee N$ have only trivial normative (descriptive) content. $D \vee N$ is compatible with all normative (descriptive) cells. Thus, $D \vee N$'s normative (descriptive) content comprises all normative (descriptive) cells which on Oddie's account means that the normative (descriptive) content is trivial. Against this taxonomic background, I now turn to one of Prior's examples. I argue that Oddie's account yields an unsatisfying assessment of this example. The reason is that his analysis cannot account for disjunctive normative content parts of mixed disjunctions. Let me first remind the reader of Prior's example: $D \vee N$, $\neg D$ entail N . N is supposed to have non-trivial normative content. The example matters to the autonomy debate because it might seem as if descriptive propositions entail a normative proposition. Let us see how to respond to this example based on Oddie's taxonomy. According to the

taxonomy, $\neg D$ is descriptive, while $D \vee N$ is neither descriptive nor normative. So, on Oddie's account one can reject the example because one of its premises is not descriptive. However, this result is not fully satisfying. To see why, note that it is still an example where premises that do not have non-trivial normative content entail a genuinely normative conclusion. This result still imposes a threat on the autonomy of normativity. Note also that it is not merely the aim to immunize an autonomy thesis against putative counterexamples. The aim is to understand how spurious counterexamples work.

The intuitive assessment of the example is relatively straightforward: one starts with two options one of which is normative and the other descriptive. For example, either today is Sunday or we ought to go to work. If one learns that one can exclude the descriptive option—because today is Friday—then one concludes that the normative option is the case, namely that we ought to go to work. A satisfying explanation of the example must be able to explain where the normative content comes from. Oddie's account commits us to say that $D \vee N$ has only trivial normative content. In other words, it cannot explain where the non-trivial normative content comes from. Thus, the account can neither explain the crucial role of the disjunctive premise nor does it provide a convincing story about why autonomists should not be worried about the example. A plausible diagnosis for why world-based accounts cannot describe the example correctly is based on their coarse-grainedness. Partitions on possible worlds such as *NORMATIVITY* do not distinguish between $D \wedge \neg N$ -worlds and $\neg D \wedge \neg N$ -worlds because they are alike in their normative aspects. Since $D \vee N$ is compatible with $D \wedge \neg N$ -worlds, it is compatible with $\neg N$ -cells. But of course, it is also compatible with N -cells. Consequently, $D \vee N$ does not count as ruling out normative possibilities which means that it is not normative. This result, however, cannot account for the fact that $D \vee N$ rules out *some* $\neg N$ -possibilities, namely $\neg D \wedge \neg N$ -worlds.

I conclude that a more fine-grained semantics is required. Some propositions seem to be related to the normative subject matter, yet do not entail a normative proposition. Hyperintensional distinctions will help to account for differences between intensionally equivalent propositions and to identify disjunctive content parts. A fine-grained taxonomy helps in turn to give a satisfying analysis of examples like the previously discussed one.

Summary and Prospect

In this chapter, I have been concerned with a series of attempts to draw distinctions between the normative and the descriptive domain. I have argued that none of the discussed taxonomies provides a satisfying solution in the context of an autonomy debate.

Let me briefly recapitulate the consequences that I have argued for:

- C1: A taxonomy must strike a balance between extremely fine-grained and coarse-grained distinctions.
- C2: Basing an autonomy thesis on a taxonomy so that the autonomy thesis is not about mixed propositions does not provide a satisfying solution to issues raised in the autonomy debate.
- C3: Taxonomic status is essential to propositions.
- C4: Normative items in our semantics are relevant to normative propositions; they are not in the same way relevant to descriptive propositions.
- C5: Normative items in our semantics are not in general irrelevant to descriptive propositions, in particular, they might not be modally independent from descriptive components.
- C6: The employed notion of relevance is a hyperintensional notion that cannot be captured in purely modal terms.
- C7: The normative-descriptive distinction requires hyperintensional semantics.

In [Chapter 3](#), I develop a taxonomy that responds to these consequences. The taxonomy is based on states rather than possible worlds. In contrast to worlds, states are not required to be complete. Roughly, they are smaller chunks of information. This will allow me to stay neutral with respect to modal relations between the normative and descriptive domain, while assuming that normative components in our semantics are relevant only to normative propositions. The account will yield a fruitful analysis of normative content that provides the ground for an autonomy debate.

Notes

- 1 Accounts assuming this taxonomy can be found in [Jackson \(1974\)](#), [Pigden \(1989\)](#), and [Schurz \(1997\)](#).
- 2 See [Brown \(2014, p. 56\)](#).
- 3 Indeed, this is what I argue for in Section ‘Intensional Taxonomies’.
- 4 Another example might be [The concept expressed by ‘ought’ is quite relevant for the ethicist].
- 5 One might argue that in fact there is only a single proposition in this case, namely the normative one. This response depends on how fine-grained propositions are individuated. Note that this result is in line with my overall argument. A satisfying taxonomy must strike a balance between opposed granularities.
- 6 See, e.g., [Brown \(2014, pp. 62–63\)](#). Brown’s account concerns ethical sentences, but it can be adjusted such that it applies to propositions.
- 7 [Maitzen \(1998\)](#) proposes a variant of this approach which states the following necessary condition for a sentence to be ethical: a sentence is ethical only if there exists a *given* moral property, *P*, that is such that for any possible

world, w , at which the sentence is true, P is instantiated at w . This restriction entails that disjunctions with ethical disjuncts involving different ethical concepts are not ethical because they do not entail that a *particular* moral property is instantiated. See Hill (2008, pp. 554, 560–562) for this construal and helpful discussion of Maitzen’s proposal.

- 8 This is a modified version of Brown’s definition of *positively ethical* sentences (Brown 2014, p. 62).
- 9 *Proof.* Assume that P entails Q , where Q is a normative proposition. I show that P is normative. According to the definition, Q entails that the extension of some normative concept is not empty. By transitivity, P entails that the extension of some normative concept is not empty. Thus, P is normative.
- 10 I assume that there are normativity-free worlds and that any proposition is true or false in a world but not both.
- 11 See Hill (2008), Pigden (2010b), Clipsham (2015), and references therein.
- 12 Consequently, the taxonomy violates the *Orthodoxy*-desideratum.
- 13 Note also that all necessary falsehoods are normative according to NOC.
- 14 Notice that the account only licenses a *collective*, in contrast to a *distributive*, reading of SLAT. On a distributive reading, SLAT entails that at least one of the premises is normative, while on a collective reading, it entails that the premises together are normative (Brown 2015). I will eventually defend a distributive interpretation of the autonomy thesis that also seems to be the more plausible reading.
- 15 For doubts see Schurz (1997, p. 11): ‘Now, to escape from Prior’s paradox by excluding mixed sentences is not a truly satisfying solution, but only a way of “defining away” the problem’.
- 16 The first occurrence of the necessity operator is commonly taken to refer to conceptual necessity; the second occurrence is more controversial. For a survey over the debate and the employed formalization, see McPherson (2021).
- 17 This is sometimes called a Jackson-style argument. Jackson (1998, ch. 5) provides an argument along these lines in favor of a conceptual analysis of ethical concepts.
- 18 One might also think that in normativity-free worlds even some simple descriptive propositions are necessarily false, e.g., that there are agents with a free will.
- 19 Russell and Restall (2010) phrase their approach in terms of sentences. However, the approach is eventually concerned with propositions. Propositions are assigned to sentences as semantic values and the inferential relation is a function of the semantic values of the sentences. Note that what I call *stability* is called *preservation* on their account. However, the stability-fragility-contrast seems more natural.
- 20 See Russell and Restall (2010, p. 247).
- 21 Recently, Russell (2021) proposes a new modified version of the taxonomy proposed in Russell and Restall (2010). Basically, the difference is that normative propositions are no longer defined in terms of R -fragility, but in terms of R -breakability. R -fragility requires that for *all* models that satisfy P there is an R -related model that falsifies P . R -breakability only requires that this is the case for some model. As a result, and in contrast to the account at issue, disjunctions and mixed material conditionals are normative. The modified account is not subject to most of the criticism raised here. However, it still cannot account for normative necessary truths because it is based on intensional semantics. In Section ‘Intensional Taxonomies’, I argue against taxonomies

- based on intensional semantics. Moreover, the new taxonomy requires a significantly different strategy concerning an autonomy thesis which will be addressed in [Chapter 4](#).
- 22 [Russell and Restall \(2010, p. 248\)](#).
 - 23 See [Russell and Restall \(2010, sec. 5.4\)](#).
 - 24 Note that being translation-stable is stronger than being not translation-fragile.
 - 25 The former follows from the definition together with the fact that if P entails Q , then every model that satisfies P satisfies Q . A counterexample to closure under negation are mixed conjunctions. They are normative according to the taxonomy, but their negation is equivalent to a mixed disjunction which is not classified as normative.
 - 26 See [Russell and Restall \(2010, p. 256\)](#).
 - 27 See, e.g., [Vranas \(2010\)](#), [Schurz \(2010\)](#), [Singer \(2015\)](#), [Maguire and Woods \(2017\)](#), and [Wolf \(2020\)](#).
 - 28 A similar assessment can be found in [Wolf \(2020: 91\)](#). However, I think, Wolf's way of presenting the issue in terms of 'uncomfortable circularity' is not as clear as it could be.
 - 29 See [Vranas \(2010\)](#) and [Schurz \(2010\)](#).
 - 30 See [Maguire and Woods \(2017, pp. 433–434\)](#).
 - 31 See [Vranas \(2010, p. 264\)](#).
 - 32 Singer refers to sentences but assumes a semantics originating from [Gibbard \(2003\)](#) that is based on possible worlds and normative standards. Thus, the account straightforwardly applies to propositions.
 - 33 [Maguire and Woods \(2017\)](#) point out that it is unclear whether the world-norm pairs model *conceptual* or *metaphysical* possibility.
 - 34 [Singer \(2015: 202\)](#).
 - 35 Given that one and only one of P and $\neg P$ is true, the proof is trivial. Let me show that the descriptive domain is closed under negation, if it is restricted to contingent propositions. Let P be contingent and descriptive. Assume for *reductio* that $\neg P$ is norm-variant. By definition, there is a w such that $\neg P$ is true in w under n but false under n' for some norms. Let $\langle w, n \rangle$ make $\neg P$ true and $\langle w, n' \rangle$ make $\neg P$ false. Thus, $\langle w, n' \rangle \models P$. Since P is norm-invariant, w makes P true under every choice of norms. Hence, $\langle w, n \rangle \models P$. This contradicts the assumption that $\langle w, n \rangle \models \neg P$. Thus, $\neg P$ is norm-invariant.
 - 36 I scrutinize Singer's explication of the autonomy thesis in [Chapter 4](#).
 - 37 See [Maguire and Woods \(2017: 426\)](#) for the reverse case.
 - 38 For the sake of the argument, I ignore that the inference is not valid.
 - 39 See [Maguire and Woods \(2017: 428\)](#).
 - 40 [Singer \(2015, pp. 208–209\)](#).
 - 41 See [Singer \(2015, p. 201, fn. 20\)](#).
 - 42 See, e.g., [Lewis \(1988\)](#).
 - 43 See [Humberstone \(1996: 151\)](#).
 - 44 This objection is made in [Maguire and Woods \(2017: 439\)](#).
 - 45 See, e.g., [Lange \(2018, p. 180\)](#): 'Had it been morally permissible to "torture" cats for fun, then cats would have had to have been radically different from the way they actually are: they would have had to have been robots or something like that, not animals'. Lange is concerned with metaphysical necessity.
 - 46 See, e.g., [Berto and Jago \(2018\)](#).
 - 47 One might wonder whether one could still adopt the taxonomy and reject Singer's autonomy thesis. As far as I can see, this is a viable option. Note that the corresponding semantics entails hyperintensional distinctions among

- propositions for which I will argue at the end of this chapter. Yet, at this stage, Singer would owe us a fully developed account.
- 48 Here is why: let D be norm-invariant and N norm-variant. Suppose that all possible world-norm pairs that make D true make N true and that some impossible world-norm pairs that make D true make N false. Thus, D entails N . However, N is norm-variant with respect to world-norm pairs that make the premises true.
- 49 This does not exclude that *distinct* taxonomies assign one and the same proposition to different categories.
- 50 I use ‘ $\llcorner \dots \lrcorner$ ’ to refer to arguments.
- 51 See [Brown \(2014\)](#), pp. 55–56.
- 52 To illustrate: Socrates does not depend on his being pale, yet he is pale.
- 53 See, e.g., [Maguire \(2015\)](#).
- 54 See, e.g., [Karmo \(1988\)](#).
- 55 See [Yablo \(2014\)](#), p. 42.
- 56 See [Brown \(2014\)](#), pp. 55–56.
- 57 See, e.g., [Lewis \(1988\)](#), [Humberstone \(1982, 1996\)](#), and [Oddie \(2018\)](#).
- 58 See [Oddie \(2018\)](#), p. 615).
- 59 Oddie takes closure to be a desirable feature in the context of an autonomy thesis ([Oddie 2018](#), p. 612). Though, see fn. 27 on the same page.

3 State-Based Taxonomy¹

The Framework: States and Exact Truthmaking

The taxonomy elaborated in this chapter is based on truthmaker semantics as developed by Kit Fine in a recent series of publications (2017a; 2017b; 2017c) and its application to the normative domain (2021). Most importantly, truthmaker semantics focuses on those parts of states of affairs that are wholly relevant to the truth of a proposition. So, one splits up possible worlds into smaller parts to be able to represent specific, wholly relevant ways for a proposition to be true. These smaller parts are what I shall call *states*. Only those states that are wholly relevant to the truth of a proposition are taken to be *exact truthmakers* of the proposition in question. The aim of this section is to introduce the formal account of states, their interrelations, and the relation of exact truthmaking.

To get a grip on the account, consider the proposition that it is sunny. In possible world semantics, this proposition is identified with the set of possible worlds in which it is sunny. These worlds make the proposition true. Possible worlds are required to be complete. That is, they settle every question. For example, all of them settle the question of whether there is a pope. However, this question is plausibly irrelevant to the truth of the proposition that it is sunny. Rather than considering whole possible worlds, truthmaker semantics thus focus on *states*. Only those parts of worlds are *exact truthmakers* of a proposition *P* that satisfy the following requirements: (i) they are rich enough to bring about *P*'s truth, (ii) they are wholly relevant to *P*'s truth. For example, the presence of sun brings about the truth of the proposition that it is sunny and is wholly relevant to it. In contrast, the presence of sun and a pope contains parts that are irrelevant to the truth of the proposition that it is sunny. Thus, the state fails to be an *exact* truthmaker of the proposition in question.

To put it in general terms: states can be fragments or proper parts of possible worlds. They can be incomplete, i.e., they need not settle every question. We remain neutral on the exact ontological status of states but

plausible candidates are facts, events, conditions, etc.² A truthmaker of a proposition must bring about and be wholly relevant to the proposition's truth. This means that if we add an arbitrary state to a truthmaker of a proposition, the result is not guaranteed to be a truthmaker of the proposition in question.

States are required to be relatively specific. Consider the proposition that it is rainy or sunny and its truthmakers. One truthmaker is the presence of sun, another is the presence of rain, but there are not in addition to that any disjunctive truthmakers. There is no state to the effect that there is rain or sun that makes the disjunction true.³ Another important feature of states is that they stand in mereological relationships to each other. For instance, the state that Socrates exists is part of the more complex state that Socrates and Plato exist. I will denote the parthood relation among states by \sqsubseteq .

I shall assume that any states have a fusion. The fusion of states s_1, s_2, \dots is supposed to be the smallest state that has all of s_1, s_2, \dots as parts. I denote the fact that s is the fusion of s_1, s_2, \dots as follows: $s = s_1 \sqcup s_2 \sqcup \dots$. For concreteness, consider a state to the effect that Socrates and Plato exist. This state contains the state that Socrates exists and the state that Plato exists and is part of every state containing these two states. The two limiting cases of fusions are the fusion of all and the fusion of zero states. I shall call the fusion of all states the full-state (\blacksquare), and the fusion of zero states, i.e., the empty fusion, the null-state (\square). The null state is part of every state.⁴

Note that not all states obtain. Suppose, for example, that Ann is 30. In this case, a state to the effect that Ann is 40 does not obtain. Moreover, some states cannot obtain. The fusion of states that cannot obtain together is an *inconsistent* state.⁵ Take, for example, a state to the effect that Plato is present and absent. This state cannot obtain. However, states need not obtain and thus the framework can allow for inconsistent states. I shall say that two states are *incompatible* iff their fusion is inconsistent and compatible otherwise. Note that possible worlds can be recovered by the fusion operation. A possible world w is a maximally consistent fusion of states. That is, w is a consistent state and the fusion of w and states that are not already part of w is always inconsistent.

A *state space* is the set of all states, \mathcal{S} , together with the parthood relation, \sqsubseteq . The parthood relation is assumed to be reflexive, anti-symmetric, and transitive.⁶ Thus, \sqsubseteq is a non-strict partial order among states. This is an important feature, since it will allow me to define propositional parts.⁷ A *modalized state space* is the triple $(\mathcal{S}, \mathcal{S}^\diamond, \sqsubseteq)$, where \mathcal{S}^\diamond is the set of consistent states. I assume that \mathcal{S}^\diamond is non-empty and closed under part. That is, every consistent state contains only consistent parts.

A proposition is in turn identified with the *set of states* that exactly make the proposition true—its exact truthmakers.⁸ I will denote the fact that s is an exact truthmaker of P as follows: $s \Vdash P$. I now turn to the truthmakers of logically complex propositions such as conjunctions and disjunctions. A state v makes a conjunction $P \wedge Q$ true iff v is the fusion of states s and t that make P and Q true, respectively. A state v makes the disjunction $P \vee Q$ true iff it makes either P or Q true.⁹

Conjunction: $P \wedge Q = \{s \sqcup t: s \Vdash P \text{ and } t \Vdash Q\}$

Disjunction: $P \vee Q = \{s: s \Vdash P \text{ or } s \Vdash Q\}$.¹⁰

The semantic clauses for conjunction and disjunction entail a distinction among necessary truths and falsities. Consider, for example, $P \vee \neg P$ and $Q \vee \neg Q$. While $s \Vdash P$ is a truthmaker of $P \vee \neg P$ it is not in general a truthmaker of $Q \vee \neg Q$. Similar reasoning applies to necessary falsities such as $P \wedge \neg P$ and $Q \wedge \neg Q$.

In a last step, I turn to negations. Consider, for example, the proposition that it is sunny (P) and the proposition that it is not the case that it is sunny ($\neg P$). The presence of sun makes the former proposition true, but is also relevant to the latter proposition. The presence of sun makes the proposition [It is not the case that it is sunny] *false*. The converse holds for states that make the negation true. The absence of sun makes [It is not the case that it is sunny] true. It makes [It is sunny] false. To account for negation, I distinguish between what I shall call a proposition's *positive* and its *negative* content. So, I extend the *unilateral* account of propositions according to which a proposition is identified with its truthmakers. On a *bilateral* account, a proposition is identified with a tuple consisting of the set of its truthmakers and the set of its *falsemakers*: $\mathbf{P} = (P^+ \ P^-)$. Falsemakers are required to be exact in the same way as truthmakers are. This allows one to account for negations. To obtain the negation of a proposition, the proposition's positive and negative content are flipped: if $\mathbf{P} = (P^+ \ P^-)$, then $\neg\mathbf{P} = (P^-, \ P^+)$. This completes the outline of state-based semantics. In the next section, I will use the formal means introduced so far to develop a state-based taxonomy of propositions.

The Descriptive-Normative Distinction

In this section, I develop a formally precise account of descriptive and normative content. The application of the state-based framework to the normative domain largely follows, but extends the taxonomy developed in [Fine \(2021\)](#). I introduce a distinction among descriptive and normative states ('Descriptive and Normative States'). Based on this distinction, I turn to propositions and define

successively less restrictive notions of normative content ('Descriptive and Normative Propositions'). Lastly, I introduce two notions of partial normative content and the notion of realized normative content ('Propositional Parts and Realized Content').

Descriptive and Normative States

In the truthmaker framework a proposition is a set of states.¹¹ To classify propositions as descriptive or normative, we need to distinguish between descriptive and normative states. Fine (2021) assumes that this distinction is primitive. Since one might take issue with this assumption, let me first address worries concerning a primitive distinction among states.

First, one might object that it is illegitimate to assume a primitive descriptive-normative distinction in a context where the descriptive-normative distinction among propositions is at issue. It is important to acknowledge, however, that the descriptive-normative distinction among propositions is motivated by appeal to an intuitive difference between paradigm cases of normative and descriptive propositions, respectively. The taxonomic project turns out to be intriguingly difficult in light of non-paradigmatic complex cases such as mixed disjunctions. Intuitively, these difficulties arise because complex cases share features of normative and of descriptive propositions. In particular, there seem to be propositions that are sometimes true in virtue of their descriptive and sometimes in virtue of their normative aspects. Let me remind the reader of essential features of states. States are assumed to be specific, i.e., non-disjunctive, and subject to part-whole relations. The outlined difficulties are significantly reduced when we turn to states. States can only combine descriptive and normative components when normative states have descriptive parts or vice versa. Thus, states allow one to first tackle a reduced problem, namely classifying mixed states. The results will in turn help to address more complex cases on the propositional level. Note that one also can adopt a companion of guilt strategy: every existing taxonomy assumes some normative primitive, e.g., normative expressions or ideal worlds.

Second, one might be worried about the ontological status of normative states. Normative truthmakers seem to require realist commitments concerning normative facts. However, we need not regard the assumption of normative states as metaphysically significant as realism about normative facts does. Following a common assumption in the debate, I assume that truthmakers in our semantics can be treated separately from metaphysical assumptions.¹² Normative states are available for anti-realists, as long as they are willing to adopt a minimalist account of normative propositions and truth. Note, however, that the account allows for a metaphysically significant interpretation. Accordingly, the framework can also serve as a basis for a metaphysical debate.¹³

With these issues out of the way, I assume that there are normative and descriptive states. A paradigm normative state is the presence of Mary's obligation. A paradigm descriptive state is the presence of sun. For the purposes of our semantics, I assume that every state is either descriptive or normative but not both. Thus, the set of descriptive states \mathcal{S}^D together with the set of normative states \mathcal{S}^N exhaust the set of all states.

A satisfying account of normative content can do justice to reasonable assumptions about the normative and the descriptive domain. It is widely recognized that some normative concepts are not purely normative. Consider, for example, thick normative concepts like *being selfish* that involve descriptive application conditions.¹⁴ Suppose that *being selfish* can be analyzed in terms of *giving priority to oneself over others plus doing it in morally bad way*. If the analysis is correct, then descriptive components figure in the normative concept. To account for this observation, I assume that normative states must contain a normative proper or improper part, yet might contain descriptive parts. In contrast, descriptive states are assumed to be *purely* descriptive. This assumption is plausible, since views according to which reality can be fully captured in descriptive terms lose their bite, if the descriptive domain includes normative components. I assume that having a normative part is sufficient to count as normative state. Formally, I assume:¹⁵

Downwards- N -Closure: If the non-empty fusion $s_1 \sqcup s_2 \sqcup \dots \in \mathcal{S}^N$, then some $s_i \in \mathcal{S}^N$.¹⁶

Upwards- N -Closure: If some state $s \in \mathcal{S}^N$ and $s \sqsubseteq t$, then $t \in \mathcal{S}^N$.

Normativity is thus a dominant feature of states.¹⁷ To avoid a trivial distinction, the null state must be descriptive, for the null-state is part of every state by assumption. Accordingly, if the null-state were normative, every state would be normative by Upwards- N -Closure.

The parthood relation among states allows us to make fine-grained distinctions between at least two types of normative states: *strictly* and *broadly* normative states. Compare the following two normative states: a state to the effect that Sarah has an obligation and a state to the effect that Sarah has an obligation and it is sunny. These states differ with respect to their descriptive parts. The latter state contains a descriptive part to the effect that it is sunny. This descriptive part is *merely added* to the normative part of the complex state. In contrast, the former state, and thereby also the latter, contains a descriptive part to the effect that Sarah exists. This descriptive part seems to be an integral component of the normative state. This suggests a distinction between *separable* and *integral* descriptive parts of normative states. A normative state n that contains a *separable* descriptive part d can be separated into the descriptive part d and a normative

state n' such that $n = d \sqcup n'$ and $n \neq n'$.¹⁸ On the other hand, a normative state n contains an *integral* descriptive part d iff (i) d is a part of n and (ii) there is no proper part n' of n such that $n = d \sqcup n'$.¹⁹ Thus, the normative state contains a descriptive part that cannot be supplemented by a normative proper part to obtain the initial normative state.

With this distinction at hand I can define strictly and broadly normative states: strictly normative states contain only integral descriptive parts; broadly normative states contain separable descriptive parts. I stay neutral with respect to the question of whether there are purely normative states that do not contain any non-trivial descriptive parts. Let me now turn to propositions.

Descriptive and Normative Propositions

I am now in a position to formulate a state-based taxonomy at the level of propositions. Basically, all proposition that can be true in normative ways are normative, while descriptive propositions can be true only in descriptive ways. The account entails an exclusive and exhaustive distinction between normative and descriptive propositions. After introducing the basic categories, I define two successively stricter normative subcategories.

Descriptive Proposition: A proposition P is descriptive iff it has only descriptive truth- and falsemakers.²⁰

Normative Proposition: A proposition P is normative iff it has at least one normative truth- or falsemaker.²¹

At this stage, I exclude the theoretical possibility that a proposition has a normative truthmaker but no normative falsemaker or vice versa.²² I shall call these propositions *regular propositions*.²³ In cases of regular propositions, it will often suffice to focus on the proposition's positive content. If not stated otherwise, I consider regular propositions. To illustrate the taxonomic results, consider the proposition [Sarah ought to help her sister]. There is a normative state that makes this proposition true, namely a state to the effect that Sarah ought to help her sister. Thus, the proposition is normative. On the other hand, consider the proposition [Sarah is in Berlin]. This proposition can be true only in descriptive ways that concern Sarah's whereabouts. Thus, it is descriptive. These examples illustrate that the taxonomy entails the intuitively correct results with respect to atomic, paradigmatic cases which means that the taxonomy meets the *Correctness-desideratum*.

I now turn to more complex cases. The taxonomy together with the semantic clauses for conjunction and disjunction entail that conjunctions

and disjunctions of two normative (descriptive) propositions are normative (descriptive). To see why, note that the fusion of normative (descriptive) states is normative (descriptive) and set-union preserves the presence (absence) of normative truthmakers. The taxonomy thus satisfies the \wedge / \vee -Closure-desideratum. Let me now turn to the assessment of mixed cases:

- (Mixed Con) Sarah ought to help her sister *and* she is in Berlin ($N \wedge D$).
 (Mixed Dis) Sarah ought to help her sister *or* she is in Berlin ($N \vee D$).

States that make the conjunction $N \wedge D$ true are fusions of normative states ($n \Vdash N$) and descriptive states ($d \Vdash D$). Since normativity is dominant, these states are normative. Hence, mixed conjunctions are normative. States that make the disjunction $N \vee D$ true are either descriptive states that make the descriptive disjunct true ($d \Vdash D$) or normative states that make the normative disjunct true ($n \Vdash N$). Again, the proposition is normative because it has at least one normative truthmaker.

One might object that disjunctions should not count as normative, for there is a purely descriptive way for them to be true. In response, I would like to stress an argument I raised against Oddie's (2018) account in the previous chapter. To assume that mixed disjunctions do not have normative content renders arguments such as $\llcorner \neg D, N \vee D \models N \gg$ unnecessarily obscure. The state-based taxonomy counts mixed disjunctions as normative. Given that we have already recognized that normative content can contain descriptive aspects, this result should not worry one too much. As will become clear in due course, the state-based account allows one to provide a detailed analysis of normative content so that one can accommodate for the intuitive difference between conjunctive and disjunctive cases. For now, it is important to acknowledge that mixed disjunctions are in the scope of an autonomy thesis because they are normative.

That said, it should be clear that the outlined notion of normative content is relatively weak. Stricter notions will allow me to take into account the uneasiness one could have when counting mixed disjunctions as normative. For this purpose, I define the categories of strictly, strongly, and weakly normative propositions.

- Strictly Normative Proposition:** A proposition P is strictly normative iff it has only strictly normative truthmakers.
Strongly Normative Proposition: A proposition P is strongly normative iff it has only strictly or broadly normative truthmakers.
Weakly Normative Proposition: A proposition P is weakly normative iff it has at least one descriptive and one normative truthmaker.




Figure 3.1 Normative categories.

Mixed conjunctions are strongly but not necessarily strictly normative, since they can have broadly normative truthmakers. Mixed disjunctions are weakly but not strongly normative, since they have descriptive truthmakers. Consequently, we have three successively less restrictive domains such that each domain is contained in its successor: $N^T \subset N^\otimes \subset N$, where N^T denotes the set of strictly, N^\otimes denotes the set of strongly, and N the set of normative propositions. By *conjoining* descriptive content, we extend the strictly normative domain and obtain the strongly normative domain. By *disjoining* descriptive content, we extend the strongly normative domain and obtain the normative domain. Figure 3.1 illustrates these relations where the rectangle represents the set of all propositions and the light gray area represents the weakly normative propositions.

Before I turn to the classification of propositions that are hard to assess, yet relevant in the autonomy debate, I introduce three further concepts that will help me to make intuitions concerning normative content formally precise.

Propositional Parts and Realized Content

We already saw that mixed disjunctions and mixed conjunctions intuitively differ. The notion of a conjunctive and a disjunctive content part will help to make this intuitive assessment precise. Additionally, the notion of realized content will help to account for the apparent context sensitivity of taxonomic status.

First, imagine someone who claims that Sarah ought to help her sister and she is in Berlin. Part of what the person claims is that Sarah ought to help her sister. Accordingly, part of what is said is normative. In contrast, the disjunction that Sarah ought to help her sister or she is in Berlin does

not commit one to either disjunct. Within truthmaker semantics we can accommodate for the intuitive difference as follows:²⁴

Conjunctive Part: P is a *conjunctive* content part of Q iff

- (i) every truthmaker of P is contained in a truthmaker of Q , and
- (ii) every truthmaker of Q contains a truthmaker of P .

Disjunctive Part: P is a *disjunctive* content part of Q iff $P \subseteq Q$.

Let $N = \{n\}$ be a normative proposition and $D = \{d\}$ a descriptive proposition. The semantic clauses for conjunction and disjunction entail: $D \wedge N = \{d \sqcup n\}$ and $D \vee N = \{d, n\}$. N is a *conjunctive* part of $D \wedge N$ because (i) n is contained in the sole truthmaker of the conjunction, and (ii) every truthmaker of the conjunction contains N 's truthmaker n . By contrast, some truthmakers of the disjunction do not contain a truthmaker of N ; for example, $d \Vdash D \vee N$ does not contain n . Hence, (ii) is violated and N is no *conjunctive* part of $D \vee N$. However, it is a *disjunctive* part since every state that makes N true also makes $D \vee N$ true. The distinction between conjunctive and disjunctive parts will become important when analyzing putative counterexamples to autonomy theses.

Since a weak notion of normative content is assumed, a normative proposition can be true in descriptive ways. The notion of the *realized content* of a proposition will turn out to be theoretically useful to account for features of mixed disjunctions. The realized content of a proposition is its overlap with reality.²⁵ Take, for instance, the proposition [Socrates is an American philosopher]. The content that overlaps with reality is that Socrates is a philosopher. To make this intuition precise, we define the greatest common part of two states. The greatest common part of two states s and t ($s \sqcap t$) is a state that is contained in both states and contains all other states that are contained in both states.

With this notion at hand, one can restrict a proposition P to particular contents by considering the greatest common parts of each truthmaker of P and a state s . Let w be the actual world, i.e., the fusion of all states obtaining in the actual world. For $P = \{p_1, p_2, \dots\}$, I shall say that $P|_w = \{p_1 \sqcap w, p_2 \sqcap w, \dots\}$ is P 's realized content. This tool helps to see why in a situation where it is true that Sarah ought to help her sister or she is in Berlin only in virtue of the fact that Sarah is in Berlin the proposition appears to be descriptive. The disjunction's realized content that Sarah is in Berlin is descriptive.

This concludes the outline of the state-based taxonomy. We are now in a position to analyze, for example, mixed disjunctions as follows: $D \vee N$ is weakly normative, yet neither strongly nor strictly normative. $D \vee N$ has

a strictly normative disjunctive part N . Lastly, $D \vee N$ can have descriptive or normative realized content.²⁶

In the remainder of this chapter, I examine classifications of propositions that are particularly hard to assess, but that play a crucial role in the autonomy debate. My results contribute to a comprehensive account of autonomy on two fronts: first, they uncover misclassifications of propositions; second, they reveal that many normative propositions that play a role in the autonomy debate are weakly normative. In the remaining two parts of this book, I provide an analysis of the threat that weakly normative propositions pose to simple interpretations of the autonomy thesis. It is important to keep in mind that the state-based taxonomy determines the taxonomic status of all propositions, but any concrete results depend on which states we consider to be truthmakers of the propositions in question. Since the semantics for several kinds of propositions are controversial or at an early stage, some results will be relative to particular semantic assumptions. However, to elucidate how we might respond to particular cases will improve our understanding of putative counterexamples to autonomy theses. I shall consider negations, conditionals, necessary truths, and propositions involving factive operators.

Classifying Negations

In this section, I consider negations of normative propositions such as [It is not permitted for Sarah to lie] (*normative negations* for short).²⁷ Normative negations are surprisingly hard to assess. On the one hand, they seem uncontroversially normative in ordinary normative thought and talk. On the other hand, to reject normativity wholesale, like the normative nihilist does, seems to require negating all normative propositions. However, this is contradictory if both N and $\neg N$ are normative. Moreover, the nihilist's insistence on the falsity of normative propositions does not seem to involve any commitment to normativity.²⁸ Let me emphasize why the classification of normative negations matters to the autonomy debate. Suppose an atomic normative proposition (N) has a descriptive necessary condition (D). It follows that the plausibly descriptive proposition $\neg D$ entails $\neg N$. Whether we face a counterexample, however, depends on whether $\neg N$ is normative. For example, if we accept the *Ought-implies-Can Principle*, it follows that [Sarah cannot help] entails [It is not the case that Sarah ought to help].

It is important to acknowledge that the state-based taxonomy entails definite results with respect to negations. Negation preserves taxonomic status of normative and descriptive regular propositions. Thus, if we consider regular propositions, the two categories are closed under negation. In other words, flipping the positive and negative content of a regular

proposition does not affect the presence or absence of normative truth- and falsemakers.

Note, however, that this result does not hold when we turn to the category of strictly or strongly normative propositions. First, consider a paradigmatic case of a merely strongly normative proposition. The mixed conjunction $D \wedge N$ is merely strongly normative given that D is not already a conjunctive part of N . Its negation $\neg(D \wedge N)$ has normative and descriptive truthmakers. On the one hand, descriptive states that make D false make $\neg(D \wedge N)$ true. On the other hand, normative states that make N false make $\neg(D \wedge N)$ true. Thus, negations of merely strongly normative propositions are weakly normative. Similar considerations apply to strictly normative propositions. Take, for example, the previously discussed case where states that make true the strictly normative proposition [Sarah ought to help her friend] contain Sarah's existence as descriptive integral part. Sarah's non-existence brings about the falsity of the proposition [Sarah ought to help her friend] and is wholly relevant to it. Therefore, it is a plausible assumption that Sarah's non-existence is a falsemaker of this proposition. Hence, it is a truthmaker of its negation. It is also plausible that the proposition [It is not the case that Sarah ought to help her friend] has normative truthmakers, e.g., the relevant permissions. Again, the negation is weakly normative. Note that in cases of strictly normative propositions, the results depend on whether one allows for purely normative states. The negation of a proposition that has only purely normative truthmakers might turn out to be strictly normative.

In a last step, I consider negations of weakly normative propositions. States that make weakly normative propositions false rule out its normative and its descriptive truthmakers. Fusions of normative states that rule out the normative truthmakers and descriptive states that rule out the descriptive truthmakers are normative. Thus, there are always normative states that make the weakly normative proposition false. However, if some normative n that makes the weakly normative proposition true contains descriptive parts, there might be also descriptive falsemakers. At any rate, the negation of a weakly normative proposition has normative truthmakers, but it might also have descriptive truthmakers. Thus, it is at least weakly normative.

The previous considerations explain why normative nihilists can accept normative negations. We already saw that normative nihilists can accept weakly normative propositions such as mixed disjunctions in virtue of accepting the descriptive disjunct. The view's skepticism is directed at strictly and strongly normative propositions. Moreover, the view entails that the realized content of weakly normative propositions is always descriptive.

One might still not find the result that negations of strictly normative propositions are classified as *weakly normative* fully satisfying. For one might object that the normativity associated with a proposition such as [It is not permitted to deceive a friend] is stronger than the normativity associated with disjunctive propositions such as [It is obligatory to help or it is rainy]. At least, it seems fine to say the following: it is obligatory to help or it is rainy, but I do not think that there is something like normativity. In contrast, it is much more confusing to say that it is not permitted to deceive a friend, but one does not think that there is something like normativity. Of course, there might be pragmatic explanations for this observation. For example, Faroldi (2014) has argued that the taxonomic status of negations depends on whether we consider internal or external negations since they make different presuppositions. He gives the following example:²⁹

(InN) Abortion is not wrong.

(ExN) It is not the case that abortion is wrong.

Faroldi argues that (InN) and (ExN) differ with respect to their normativity. He assumes that normative propositions like [Abortion is wrong] presuppose *normativity*.³⁰ This means that the proposition presupposes, roughly, that abortion can be an object of genuine moral judgment. Faroldi argues that internal negations such as (InN) keep the proposition normative, while inverting the *polarity*, e.g., from good to bad, from obligatory to forbidden, etc. By contrast, external negations such as (ExN) cancel the presupposition of normativity. As a result, the proposition need not be normative anymore which renders it acceptable to normative nihilists.

It would be interesting to see how truthmaker-based accounts of presuppositions deal with these cases.³¹ However, to introduce the details of the accounts that are still at an early stage would take us too far afield. Instead, I will show how the distinction between strictly and weakly normative propositions can account for the claim that both [Abortion is wrong] (N) and its internal negation [Abortion is not wrong] (denoted by $\sim N$) are rejected by the normative nihilist, while [It is not the case that abortion is wrong] ($\neg N$) can be accepted. Assume that n_w is a state to the effect that abortion is wrong, n_r is a state to the effect that abortion is right, and d is a descriptive state that is compatible with a normative nihilist position and represents a descriptive way for $\neg N$ to be true. If the nihilist is right and there is no normativity, d obtains and $\neg N$ is true, while both normatively binding propositions N and $\sim N$ are false. If, on the other hand, abortion is wrong, n_w obtains and N is true, while both negations $\sim N$ and $\neg N$ are false. Lastly, if abortion is right, n_r obtains and N is false, while both $\sim N$

and $\neg N$ are true. Here is the positive and negative content of the three propositions:

$$N = (N^+, N^-) = (\{n_w\}, \{n_r, d\})$$

$$\sim N = (\sim N^+, \sim N^-) = (\{n_r\}, \{n_w, d\})$$

$$\neg N = (\neg N^+, \neg N^-) = (N^-, N^+) = (\{n_r, d\}, \{n_w\})$$

Note that only $\neg N$ has a descriptive truthmaker. Thus, $\neg N$ is weakly normative. N and $\sim N$ have only strictly normative truthmakers and thus are strictly normative. In this sense, both N and $\sim N$ are normatively binding. We thereby can explain why stating that it is not permitted to deceive a friend seems to bear a stronger normative commitment than, e.g., mixed disjunctions. The internal reading can be modeled as a strictly normative proposition and thus would only be true, if the world turns out to be in a genuinely normative state.

This concludes my discussion of negations. In the autonomy debate, it is of great significance to recognize features of negations of normative propositions. Autonomy theses face putative counterexamples based on normative negations, e.g., [Sarah cannot help] entails [It is not the case that Sarah ought to help]. To defend normative autonomy existing taxonomies often exclude negations from the normative domain. This classification, however, cannot do justice to ordinary normative thought and talk. The state-based taxonomy entails the intuitively correct results. To uncover that external negations are weakly normative and thus can be true in descriptive ways provides the ground to explain putative counterexamples and why they are spurious. Moreover, the state-based taxonomy accommodates for the intuitive assessment that some negations of normative propositions carry strong normative commitments.

Classifying Conditionals

In this section, I turn to the classification of conditionals. Let me first explain why conditionals play an important role in the autonomy debate. There are two main reasons. The first reason is that autonomists claim that normative conclusions cannot be inferred just from descriptive premises; they insist on the claim that that normative conclusions can only be inferred from descriptive premises *together* with normative auxiliary premises. So-called bridge principles are required to bridge the alleged gap between the descriptive and the normative domain. It is of great significance for the plausibility of any autonomist position that these auxiliary premises or bridge

principles are classified as normative. Bridge principles, however, often have the form of conditionals with a descriptive antecedent and a normative consequent: *If D, then N*. The second reason is that some putative counterexamples to the autonomy thesis involve conditionals. The success of those examples depends on the classification of the involved propositions. Let me restate a putative counterexample that is based on a conditional. Consider the descriptive proposition that all undertakers are Church officers. This proposition entails that if all Church officers ought to be reverent, then all undertakers ought to be reverent. The conclusion is a candidate for a normative proposition. Thus, we need an accurate assessment of such conditionals to be able to evaluate the corresponding examples.

One advantage of the state-based taxonomy is that it provides us with a general criterion for the classification of propositions. The taxonomic status of a proposition is fully determined by the corresponding set of truthmakers and their classification. We do not need a definition by cases. Thus, to tell whether the relevant conditionals are normative, we just have to examine whether some normative states make the conditional true. However, as I already noted, the results depend on which states are considered to be truthmakers of the relevant conditional. There is an ongoing debate about the correct semantics for conditionals.³² I cannot settle questions concerning the correct semantics for conditionals here. To deal with this difficulty, I shall focus on two plausible interpretations of conditionals and examine the taxonomic results for both of them.

One option that I shall consider is to interpret conditionals in terms of the material conditional. Under this interpretation *If P, then Q* is true iff the antecedent is false or the consequent is true. Yet, the fact that the conditional is already true if the antecedent is false has been criticized. So, in addition, I shall consider an interpretation according to which the conditional is not already true if the antecedent is false. On that interpretation, a conditional tells us how to pass from the antecedent to the consequent. Roughly, a conditional thus understood is true iff something fills or bridges the gap between the antecedent and the consequent. Note that the metaphors that are at work in the previous description and those being exploited in the label *bridge principle* fit together. Fine (2014) discusses an interpretation of conditionals according to which a conditional tells us how to pass from the antecedent to the consequent under the label *incremental conditional*.³³ In what follows, I shall outline the taxonomic results for both interpretations of the conditional.

Let me first address the material conditional $P \rightarrow Q := \neg P \vee Q$. The state-based taxonomy entails that mixed conditionals of the form $D \rightarrow N$ or $N \rightarrow D$ are weakly normative. The material conditional is defined in terms of disjunction and negation. Thus, in both cases one has a disjunction with a descriptive and a normative disjunct. This result allows me to

apply the taxonomic results for mixed disjunctions according to which mixed disjunctions are weakly normative. Furthermore, even conditionals with a strictly normative antecedent and a strictly normative consequent $N_1 \rightarrow N_2$ are weakly normative. The previous section revealed that the (external) negation of a normative proposition—here $\neg N_1$ —is weakly normative. Since states that make $\neg N_1$ true also make $N_1 \rightarrow N_2$ true, the conditional has a descriptive truthmaker. Thus, material conditionals with a normative antecedent or a normative consequent are weakly normative. The significance of this result will become transparent later when I turn to autonomy theses.

For now, let me turn to an alternative interpretation of conditionals. The following examples motivates to consider the so-called incremental conditional:

(GUN) If your hand is holding a loaded gun aimed at a harmless person, then it is good to move your trigger finger in a shooting motion.³⁴

Intuitively, this proposition is false whether one holds a gun aimed at a harmless person or not. Read in terms of the material conditional, the proposition is true in a world where one does not hold a loaded gun at a harmless person. By contrast, the incremental conditional is not already true, if the antecedent is false.³⁵

To assess the taxonomic status of the incremental conditional $P \supset Q$, I need to clarify which states make it true. Fine (2014) provides semantics for the incremental conditional based on states telling us how to pass from a state s to a state t . These states are called conditional connections, denoted by $s \rightarrow t$.³⁶ The main idea is that a state that makes the conditional true tells us how to pass from truthmakers of the antecedent to truthmakers of the consequent. The technical details of the account would require a more lengthily introduction. For our purpose, it suffices to focus on particular features of those states.

First, conditional connections are supposed to guarantee that a state t obtains whenever a particular state s and the corresponding conditional connection obtains. Formally, this is interpreted as follows: $s \sqcup (s \rightarrow t) \sqsupseteq t$.³⁷ The condition is also satisfied by states that contain parts that are wholly irrelevant to the question of how to pass from s to t , for t is only required to be part of the fusion. To ensure that the conditional connection is wholly relevant to $P \supset Q$, it is required that $s \rightarrow t$ —which is also called the *residuation state*—is the smallest state that satisfies the former requirement. This requirement is satisfied if we define $s \rightarrow t$ as follows

Residuation state: $s \rightarrow t := \sqcap \{u: s \sqcup u \sqsupseteq t\}$

and assume that the state space is *residuated*. A state space is residuated iff $s \sqcup (s \rightarrow t) \sqsupseteq t$ holds for all states in \mathcal{S} . For definite propositions $P = \{p\}$, $Q = \{q\}$, i.e., propositions with a sole truthmaker, the truthmaker of $P \supset Q$ is $p \rightarrow q$. Additionally, there is a clause for propositions with several truthmakers based on fusions of residuation states.³⁸ However, I will mainly focus on definite propositions.

Let me illustrate that the account entails the intuitive results with respect to GUN. The incremental conditional is true iff a state obtains that tells us how to pass from a state to the effect that one is holding a loaded gun at a harmless person to a state to the effect that moving one's finger in a shooting motion is good. It is plausible that no such state actually obtains, while it might be that a state obtains that tells us how to pass from the former state to the state that moving one's trigger finger is *bad*. Thus, GUN is false.

With the truthmaker-based analysis at hand, the classification of incremental conditionals depends on the classification of residuation states. Recall that the incremental conditional is normative iff it has at least one normative truth- or falsemaker. My aim is to gain some systematic insights into particular classes of propositions such as mixed or pure conditionals. First, we scrutinize a conditional connection between a descriptive state d and a normative state n . I show that $d \rightarrow n$ is a normative state. Recall that $d \rightarrow n$ satisfies $d \sqcup (d \rightarrow n) \sqsupseteq n$. By Upwards- N -Closure $d \sqcup (d \rightarrow n)$ is normative since it contains a normative part n . By Downwards- N -Closure, we know that whenever a normative state is the fusion of other states one of the parts must be normative. Since d is descriptive by assumption, $d \rightarrow n$ is a normative state. Accordingly, propositions of the form $D \supset N$ are at least strongly normative. The difference between strongly and weakly normative propositions becomes relevant, if we consider propositions that have several truthmakers. It is easily verified that conditionals with a strongly normative consequent are strongly normative, while conditionals with a weakly normative consequent are weakly normative. The reason is that strongly normative propositions have only normative truthmakers. Thus, all states that make conditionals with a strongly normative consequent true must contain normative residuation states of the form $d \rightarrow n$. In contrast, weakly normative propositions have descriptive truthmakers. Thus, states that make conditionals with weakly normative consequent true can be fusions of descriptive residuation states of the form $d_1 \rightarrow d_2$ where d_2 is a descriptive truthmaker of the consequent. At any rate, $D \supset N$ is normative which is crucial for a coherent autonomist position because autonomists claim that we cannot infer normative conclusions from descriptive premises but need *normative* bridge principles.³⁹

I now turn to conditional connections between two normative states m and n . I show that $n \rightarrow m$ can be descriptive or normative. Thereafter,

I explain why this result is important for a correct assessment of putative counterexamples. First, I show that $n \rightarrow m$ can be descriptive. Let $m := n \sqcup d$ for descriptive d , then $n \rightarrow m \sqsubseteq d$. Since descriptive states contain only descriptive parts, $n \rightarrow m$ is descriptive. To see that $n \rightarrow m$ can be normative, assume that $m := n \sqcup n'$, where there is no descriptive state d such that $n' = n \sqcup d$. In those cases, $n \rightarrow m$ is normative. The former condition excluding that n' is the fusion of n and a descriptive state d is satisfied in many cases. Suppose, for example, that n is the state that Ann ought to help and n' is the state that Sue is permitted to run.

The previous insights help to assess particularly hard cases. Consider the conditional [If all Church officers ought to be reverent, then all undertakers ought to be reverent]. My aim is to point out that the conditional plausibly has normative and descriptive truthmakers and thus is weakly normative. Any state that makes this conditional true must tell us how to pass from states that make it true that all Church officers ought to be reverent to states that make it true that all undertakers ought to be reverent. Moreover, the assessment of this case depends on which states are considered to be truthmakers of universally quantified propositions. A plausible candidate for a state that makes the proposition [All Church officers ought to be reverent] true is a state to the effect that Sarah, Paul, and Samuel are Church officers, while there are no other Church officers, and Sarah, Paul, and Samuel ought to be reverent. Likewise, a plausible candidate for a state that makes the proposition [All undertakers ought to be reverent] true is a state to the effect that Sarah and Samuel are undertakers, while there are no other undertakers, and Sarah and Samuel ought to be reverent. Let me explain how one might pass from the former state to the latter state. Since the former state already contains the normative state that Sarah and Samuel ought to be reverent, just the descriptive part of the latter state is required, namely that Sarah and Samuel are undertakers and that there are no other undertakers. So, in this case, the conditional connection is descriptive. Note, however, that this is only one way for the antecedent to be true, while the conditional is true iff for every truthmaker of the antecedent there is a conditional connection between it and a truthmaker of the consequent. Nothing hinges on the particular choice of Church officers, though. For each truthmaker of the antecedent one can pick a truthmaker of the consequent such that the former state already contains the normative parts of the latter state. If one does so, there are only descriptive conditional connections and thus a descriptive state that makes the conditional true. This result illustrates that normativity is not essential to the example. Consider the conditional [If all Church officers wear green shoes, then all undertakers wear green shoes]. One can follow exactly the same pattern as before to obtain a state that makes this conditional true. In fact, the two conditionals share at least one truthmaker.

In contrast, the conditional connection between (i) the state that Sarah, Paul, and Samuel are Church officers, while there are no other Church officers, and Sarah, Paul, and Samuel ought to be reverent and (ii) the state that Sarah and Andrea are undertakers, while there are no other undertakers, and Sarah and Andrea ought to be reverent must contain a normative part. The normative state that Andrea ought to be reverent is required to pass from the state (i) to the state (ii). However, any state that contains this normative state is normative. Consequently, the conditional has normative truthmakers and descriptive truthmakers which means that it is weakly normative.⁴⁰

Before I conclude, let me highlight a taxonomic result that might be problematic. Consider the trivial conditional: $N \supset N$, e.g., if murder is morally wrong, then murder is morally wrong. Not even the normative nihilist denies its truth. The present account is compatible with this result. If N is a definite proposition, then nothing is required to pass from a truthmaker of N to a truthmaker of N . Thus, $n \rightarrow n = \square$, which is descriptive. Yet, and this is the problematic consequence, this might also hold for normative analytic truths such as $(OP \wedge OQ) \supset OP$. Again, nothing is required to pass from $o_p \sqcup o_q$ to o_p . This suggests that some *prima facie* normative analytic truths are descriptive and moreover, collapse into a single proposition, namely $\{\square\}$. Let me hint at a putative way to respond. So far, I focused on the positive content of conditionals. The negative content of conditionals requires technically sophisticated tools.⁴¹ Plausible candidates for states that make the conditional false are fusions of states that make the antecedent true and states that make the consequent false. This provides us with a way to distinguish between, for example, $N \supset N$ and $(OP \wedge OQ) \supset OP$ in virtue of their negative contents. Accordingly, propositions might have trivial and thereby descriptive positive content and normative negative content. Such cases generate exceptions from the assumption that propositions with normative positive content have normative negative content and vice versa. Exceptions are limited to cases where either the proposition's positive or its negative content is trivial. For that reason, I had to restrict the results concerning negations to regular propositions, i.e., propositions with non-trivial positive and negative content.

Let me summarize the main results: conditionals with a normative consequent are at least weakly normative. This result provides the ground for a strategy to react to putative counterexamples. Conditionals of the form $D \supset N$ that might be considered to bridge a putative gap between the normative and the descriptive domain are often strongly normative. This result makes the autonomist position viable and provides the resources to meet the *Orthodoxy*-desideratum according to which moral laws are normative.

Classifying Necessary Truths

I shall now turn to different kinds of necessary truths. Recall that intensional accounts of propositions face serious difficulties when it comes to the classification of necessarily equivalent propositions. Part of what is at issue in the autonomy debate is whether descriptive propositions can entail normative propositions. Now, suppose we face a putative two-way counterexample. That is, an example where a seemingly normative proposition entails and is entailed by a seemingly descriptive proposition. Consider, for example, a view according to which necessarily, an act is morally good if and only if the act maximizes happiness. Since intensional accounts do not distinguish between necessarily equivalent propositions, they cannot even recognize this example as a putative counterexample to an autonomy thesis that is concerned with propositions.⁴² The intensional account must say that the one proposition is either normative or descriptive but not both given that the categories are exclusive. Note that if the categories are not exclusive, one would obtain trivial counterexamples because every descriptive and normative proposition entails itself. One might object that intensional accounts can still distinguish between the corresponding sentences, but this is exactly what a hyperintensional account will do on the level of propositions while avoiding a commitment to a particular language.

The state-based account can even distinguish between different logically necessary truths (falsities). Let N be a normative and D is a descriptive proposition. $N \vee \neg N$ and $D \vee \neg D$ are distinct. Moreover, they belong to different categories because the former but not the latter has normative truthmakers. Thus, $N \vee \neg N$ is normative while $D \vee \neg D$ is descriptive. Similar considerations hold for the outlined putative two-way counterexample.

Note that this result is controversial. For example, [Prior \(1960\)](#) claims that logical truths that involve normative components 'belong to the logic of ethics' but not to ethics.⁴³ However, there is, I think, a reason to take the taxonomic result to be more than a technical consequence of closure principles. The first thing to note is that $N \vee \neg N$ is always true in virtue of either N or $\neg N$. So, one may say that the necessary truth is rooted in its true disjunct.⁴⁴ Both disjuncts are normative according to the state-based taxonomy. So, there is or at least can be a normative reason why the logical truth holds.

But why should we care for the taxonomic status of necessary truths? It is easy to see that in classical logic, necessary truths and falsities give rise to putative counterexamples to a simple logical autonomy thesis, e.g., $D \models N \vee \neg N$. It is commonly agreed that these cases are spurious counterexamples. There are, however, different explanations why this is so. Intensional accounts often point out that there is just a single logical truth, namely the set of all possible worlds. This suggests that, if the logical truth

has any normative content, it has only trivial normative content.⁴⁵ On that basis *all* examples that involve necessary truths or falsities can be rejected. The inferred normative conclusion or the assumed descriptive premises are not genuinely normative or descriptive, respectively.⁴⁶ However, some examples suggest that a more fine-grained assessment of those cases improves our understanding of spurious counterexamples. Consider a putative genuine counterexample where contingent descriptive D entails contingent normative N . Now, reconsider the above pattern: $D \models N \vee \neg N$. This entailment relation holds not just because $N \vee \neg N$ is a necessary truth, but also because D entails a normative disjunct. To recognize the normative content of some necessary truths allows us to distinguish between the previous case and cases where the descriptive premises are unrelated to the normative content of the conclusion. It strikes me as a plausible result that the status of spurious counterexamples does not merely depend on $N \vee \neg N$'s allegedly trivial content, but on the question how the premises relate to the conclusion. Roughly, the idea is that *only if* D does not properly relate to $N \vee \neg N$'s normative content we face a spurious counterexample.⁴⁵ Similar considerations hold for necessary falsities. So, one advantage of the state-based taxonomy is that it avoids overgeneralization by drawing a more accurate picture of putative counterexamples.

In a second step, one might ask whether all logically necessary truths that can be expressed in normative terms are normative. The answer is determined by the state-based taxonomy. Whenever a logically necessary truth has at least one normative truth- or falsemaker it is normative. A concrete response, however, would require a comprehensive semantic account of all logically necessary truths which cannot be discussed here. Note that, if the previous analysis is correct, the incremental conditional $N \supset N$ is a logical truth that can be expressed in normative terms but does not have normative *positive* content.

In addition to logical truths, I shall now turn to weaker forms of necessity⁴⁷ and discuss the classification of *prima facie* normative modal propositions. Before I explain why these cases are worth being addressed, let me make an important caveat that we already touched upon. Truthmaker semantics for modal operators are at an early stage. It would go beyond the scope of this discussion to introduce the details. To deal with this difficulty, I will merely outline some features of states one might consider making the propositions in question true which allows me to stay largely neutral with respect to concrete accounts.

The new focus brings us back to the question of why we should be concerned with the taxonomic status of necessities other than logical ones. The first reason is a strategic one. Variance-based taxonomies have difficulties to categorize *prima facie* normative conceptual and maybe even metaphysical truths as normative. They have been criticized for conflicting

with our intuitive assessment of those cases. So, if the state-based taxonomy can improve on these difficulties, this counts in favor of the overall account. The second reason is more substantial: moral laws are paradigm examples of normative propositions. Recall that according to the *Orthodoxy-desideratum*, a taxonomy must classify moral laws and normative principles as normative. Commonly, it is assumed that moral laws hold with either metaphysical or normative necessity.⁴⁸ So, for a satisfying account of normative content the assessment of modal propositions is important. Moreover, as soon as we turn to metaphysical autonomy these the status of metaphysical necessities will become relevant.

Let me remind the reader of a case where the model-variance approach entails unintuitive results. With respect to sets of models that satisfy $P \rightarrow OP$ and P , OP turns out to be non-normative because there is no model \mathfrak{M} that does not satisfy OP . By contrast, the state-based taxonomy entails that OP is normative because it is based on an *exact* notion of truthmaking. The account focusses on states that are wholly relevant to OP 's truth instead of complete possible worlds. At any rate, the presumed necessity of OP provides no reason to think that an exact truthmaker of OP would have to be non-normative.

Let me now turn to moral laws or normative principles. I focus my attention on laws that specify descriptive antecedent conditions. For example, one might consider the following candidate for a moral law: necessarily, if one promised to ϕ , then one ought ϕ . In general, I consider propositions of the form: necessarily, if D , then N . I shall focus on interpretations in terms of the incremental conditional $\Box(D \supset N)$. However, the strategy to deal with the \Box -operator is equally applicable to interpretations in terms of the material conditional $\Box(\neg D \vee N)$. It is easily verified that on the outlined account $\Box(\neg D \vee N)$ is weakly normative. I shall assume that the strategy to deal with the \Box -operator can be applied to different forms of necessity, if it is adjusted accordingly.

Recall that the conditional connection $d \rightarrow n$ that makes $D \supset N$ true is a normative state. However, in general, $d \rightarrow n$ will not suffice to make $\Box(D \supset N)$ true. What is needed for $\Box(D \supset N)$ to be true in a world is that $D \supset N$ is true in all (accessible) worlds. So, roughly, a truthmaker of $\Box(D \supset N)$ might be the fusion of states that make $D \supset N$ true relative to each of the relevant worlds.⁴⁹ What matters to our purposes, however, is the weaker, independently plausible assumption that states that make $D \supset N$ true are part of states that make $\Box(D \supset N)$ true. A state that tells us how to pass from, for example, a promise to ϕ to an obligation to ϕ plausibly helps to make it true that necessarily, we ought to ϕ , whenever we promised to ϕ . Given this is correct, normative states such as $d \rightarrow n$ are part of states that make $\Box(D \supset N)$ true. Thus, $\Box(D \supset N)$ has normative truthmakers. It is important to acknowledge that given that the previous

sketch of a truthmaker-based analysis is correct, truthmakers of $\Box(D \supset N)$ where N is strongly normative are at least strongly normative. By contrast, $\Box(\neg D \vee N)$ is merely weakly normative according to the discussed account. Since $\neg D \vee N$ is merely weakly normative, truthmakers of $\neg D \vee N$ that help to make $\Box(\neg D \vee N)$ true need not be normative. In both cases, the state-based taxonomy can do justice to autonomists, according to which normative premises are required to *bridge* a gap between the normative and the descriptive domain.⁵⁰

To conclude let me emphasize two main advantages of the state-based taxonomy in its application to necessary truths: (i) a fine-grained analysis of logically necessary truths that allows us to make relevant distinctions between putative counterexamples; (ii) the classification of moral laws which specify descriptive antecedent conditions as normative.

Classifying Propositions Involving Factive Operators

In this section, I argue that propositions that can be expressed by applying a factive operator to a normative proposition are normative. An operator \mathcal{O} is factive iff whenever it is applied to a proposition P , $\mathcal{O}(P)$ entails P . In particular, I am concerned with propositions that result from applying the truth operator *it is true that* to normative N . The classification of those propositions in turn provides the ground to assess a number of counterexamples to autonomy theses from the literature.

To begin with, let me introduce a notational convention. I shall call propositions that can be expressed by applying a factive operator to a proposition *factive propositions*. One might wonder why factive propositions should be explicitly addressed. In response, let me highlight three important reasons. First, as already indicated, a number of putative counterexamples to autonomy theses involve factive propositions among the premises or a conclusion that is a factive proposition. To assess these examples, it will be illuminating to discuss systematic similarities and issues that arise with respect to those cases. Second, and relatedly, according to the *Specificity*-desideratum, a taxonomy must be fine-grained enough to account for relevant differences between structurally similar cases. As I will argue, whether propositions are factive is a criterion by means of which relevant differences can be characterized. Third, it is worth admitting head-on that some classifications concerning factive propositions might on the face of it seem counterintuitive. To mitigate these worries, it is important to acknowledge the broader context of these results.

Before elaborating on each of these points in turn, I present the taxonomic results for factive propositions involving the truth-operator τ . This paradigmatic factive operator will give a good idea of why factive propositions matter to autonomy theses. For concreteness, consider the

proposition [It is true that Sarah ought to help Ann]. To determine the taxonomic status of this proposition we must determine which states make the proposition true. Suppose that the proposition [Sarah ought to help Ann] is made true by the normative state n . This state seems to be wholly relevant to the truth of the factive proposition [It is true that Sarah ought to help Ann]. Thus, n is at least part of a truthmaker of the factive proposition which means that the proposition is at least weakly normative. In what follows, I will assume that the application of the truth operator τ to P does neither affect P 's positive nor P 's negative content. Note, however, that the result that $\tau(P)$ is normative also holds on the weaker assumption that P 's truthmakers are only parts of $\tau(P)$'s truthmakers. So, $\tau(P)$ belongs to the same categories as P given that

$$\text{Truth: } \tau(P) = \left([\tau(P)]^+, [\tau(P)]^- \right) = (P^+, P^-) = P.$$

This result is convenient in the context of autonomy theses. To see why, assume that $\tau(N)$ for normative N had been descriptive. Since $\tau(N)$ entails N there would have been counterexamples to autonomy theses for each normative proposition. Of course, this is not the justification for the taxonomic result (otherwise, the taxonomy would not be impartial). It is independently plausible that the classification is correct since τ does not seem to change N 's (exact) truth-conditions. Let us now turn to a putative counterexample from the literature where the truth operator plays an important role:⁵¹

(Truth-Teller*) Everything Alfie says is true (P1).
 Alfie says that it is impermissible to starve the Irish (P2).
Thus, it is impermissible to starve the Irish (C).

Let us first scrutinize P1. According to the state-based taxonomy, P1 is weakly normative. If the reader is already convinced that this is the correct result, she might skip this paragraph. Otherwise, it might help to consider cases in between examples such as [It is true that Sarah ought to help Ann] and P1. Consider the proposition that what Sarah just said is true and its truthmakers.⁵² Suppose the proposition is actually true. Suppose further that Sarah just said that she ought to help Ann. The actual truthmaker of the proposition seems to have two parts: (i) a state to the effect that Sarah just said that she ought to help Ann and (ii) a state to the effect that Sarah ought to help Ann. Such a state is normative due to its normative part (ii). Thus, the proposition is at least weakly normative. Now consider P1. Since Alfie might have said several things, we might think of states that make P1 true as fusions of states that are structurally similar to the state described in Sarah's case. Like before, P1's actual truthmakers

are normative iff Alfie said something normative. Thus, P1 is guaranteed to be normative. Of course, there are also descriptive ways for P1 to be true. P1's actual truthmakers are descriptive iff Alfie said only descriptive things. Thus, P1 is indeed weakly normative.

It is important to acknowledge that the validity of the argument depends on the fact that the first premise is factive. Suppose the first premise were [Everything Alfie says is interesting], then the argument would not be valid. However, not only the validity of the argument depends on whether the proposition is factive, but also its normative status. To see that, just note that non-factive operators such as *Alfie says that* can be applied to normative propositions while the resulting proposition is descriptive. For example, P2 is of the form *Alfie says that N*, but is arguably descriptive. The truthconditions of the proposition do not require the world to be in a particular normative state. They merely require that Alfie uttered something with a particular meaning.

Given the factive–non-factive distinction, it becomes transparent how the state-based taxonomy meets the *Specificity*-desideratum. I motivated the desideratum by means of an intuitive difference between the propositions [Sarah knows that she ought to help Ann] and [Sarah believes that she ought to help Ann]. It is commonly assumed that the knowledge operator is factive, while the belief operator is non-factive.⁵³ Thus, the former but not the latter proposition requires the world to be in a particular normative state. Normative states are part of states that make true propositions that attribute normative knowledge, yet not of states that make true proposition that attribute normative beliefs.

In the remainder of this section, I address some worries that one might have concerning the result that [What A says is true] is classified as weakly normative. One might object that this taxonomic result is only plausible in circumstances where A says something normative, while whenever what A says is descriptive the proposition [What A says is true] is descriptive. Yet, this taxonomic result cannot be achieved on an *essentialist* view given that it assumes that in both circumstances the same proposition is expressed. There are two points to stress in response: first, given the assumptions about propositions such as [What A says is true] all kinds of states are parts of truthmakers for that proposition. A could say everything and thus, the proposition is in a weak sense also about mathematics, hedgehogs, astronomy, etc. These observations relativize the result that the proposition is weakly normative. Yet, one might insist that instead of a single implausible taxonomic result, we now have many. The second consideration addresses this worry. Truthmaker semantics provide us with the formal tools to account for a proposition's *realized content*. Recall that P's realized content is P's overlap with reality. Thus, we can compare the realized content of [What A says is true] in circumstances C_N where A says

that it is impermissible to starve the Irish and in circumstances C_D where A says that Jupiter is the biggest planet. In C_N , the realized content of [What A says is true] is strictly normative. In C_D , it is descriptive; in particular, the realized content in C_D is in a strict sense about astronomy. Similar considerations apply to proposition such as [There is a truth] and, in particular, to the proposition [There is a normative truth]. For strategic reasons, I do not spend too much place on discussing these cases. These propositions quantify over a domain to which they belong which gives rise to well-known complications that cannot be discussed here.⁵⁴ This is the reason why I focused my attention on cases that do not involve such additional complications. This concludes my discussion of factive propositions.

Summary

The state-based taxonomy distinguishes between different ways for a proposition to be true that are wholly relevant to its truth. In particular, a distinction between descriptive and normative states that exactly make a proposition true provides the ground to make fine-grained distinctions between normative and descriptive content and content parts. Classifications on the level of states allow us to define notions of normative content that are successively less demanding concerning the purity of normative content: strictly, strongly, and weakly normative propositions. Setting aside irregular cases, the taxonomy entails that the normative and the descriptive category are closed under the Boolean operations. By contrast, neither domain is closed under converse-entailment. The account makes hyperintensional distinctions and thus avoids short-comings of intensional accounts. The taxonomy together with its formal tools sheds light on complex cases discussed in the literature. It uncovers and explains putative counterexamples to autonomy theses that are based on misclassifications like (Truth-Teller) and (Normative Knowledge). It does so by uncovering the role played by factive operators. Moreover, the taxonomy accommodates and explains competing intuitions concerning negations of normative propositions. Finally, it provides a satisfying classification of conditionals, in particular moral laws, which in turn makes room for a coherent autonomist view according to which bridge principles are normative.

Notes

- 1 Parts of this chapter are based on a previously published Article in Analysis: ‘A Semantics for Moral Error Theory’ (2024), *Analysis* 84 (2), 221–230.
- 2 At least under a realist interpretation. I return to the ontological status of states when I turn to normative states.
- 3 For a discussion of disjunctive truthmakers, see, e.g., [Jago \(2023\)](#).

- 4 To assume the existence of the null-state is required for largely technical reasons. Note that the full-state will often be an impossible state.
- 5 I use the terms ‘(in-)consistent’ and ‘(im-)possible’ interchangeably.
- 6 *Reflexivity*: $s \sqsubseteq s$. *Anti-symmetry*: if $s \sqsubseteq t$ and $t \sqsubseteq s$, then $s = t$. *Transitivity*: if $s \sqsubseteq t$ and $t \sqsubseteq u$, then $s \sqsubseteq u$.
- 7 I introduce the notion of propositional parts in Section ‘Propositional Parts and Realized Content’.
- 8 I will often drop the qualification. If not indicated otherwise, I refer to exact truthmaking.
- 9 In what follows, I shall distinguish between the positive and the negative content of a proposition. The following clauses specify the positive content of conjunctions and disjunctions, respectively.
- 10 Concerning disjunctions, one might prefer to work with the closure under fusion of $\{s : s \Vdash P \text{ or } s \Vdash Q\}$. The closure under fusion is the smallest set that includes all fusions of states in $\{s : s \Vdash P \text{ or } s \Vdash Q\}$. For the current purposes, the distinction does not matter. When I turn to autonomy theses, I shall assume that disjunctions are closed under fusion.
- 11 For simplicity’s sake, I will work with the unilateral conception of propositions as long as I am not concerned with negations.
- 12 See, e.g., [Fine \(2021\)](#) and [Jago \(2020\)](#).
- 13 Let me emphasize that the truthmaker account can serve as a *basis*, but might be rejected as a *guide* to metaphysics. See, e.g., [Schneider \(2006\)](#).
- 14 For an overview, see [Väyrynen \(2021\)](#). In [Chapter 6](#), I will return to the distinction between thick and thin normative concepts and discuss in detail how it bears on the autonomy debate. For an early version of the view that normative concepts that are commonly taken to be thin contain descriptive components, see [Bolzano \(1837: §144\)](#).
- 15 See [Fine \(2021, p. 908\)](#).
- 16 Or equivalently, given that \mathcal{S}^N and \mathcal{S}^D are exclusive and together exhaustive: if $s_1, s_2, \dots \in \mathcal{S}^D$, then $s_1 \sqcup s_2 \sqcup \dots \in \mathcal{S}^D$.
- 17 See [Fine \(2021, p. 908\)](#).
- 18 One might wonder whether we should require that the separable descriptive and normative parts do not share any parts. Yet, a state to the effect that Sarah has an obligation and prefers green shoes is plausibly broadly normative, while both the descriptive and the normative part arguably contain Sarah’s existence.
- 19 This assumption violates the mereological principle of weak supplementation according to which whenever x is a proper part of y there is a z such that z is a part of y that does not overlap x . One might be worried by this consequence. I will sketch two reasons to not be worried here. First, since states are a technical tool, one might not expect them to have the same mereological features as ordinary objects. Second, the mereological principle would require us to assume purely normative states which is itself a fairly controversial assumption in the present context. See [Fine \(2021, p. 911\)](#).
- 20 See [Fine \(2021, p. 893\)](#).
- 21 See [Fine \(2021, p. 913\)](#).
- 22 In principle, it is possible to fix positive and negative content independently. However, since we allow for inconsistent states, it seems reasonable to assume that for each normative state n there are (possibly inconsistent) normative *converse* states. That is, normative states that are wholly relevant to rule out n . Now, consider a proposition P that has at least one normative truthmaker n .

- P 's falsemakers must be incompatible with all its truthmakers, and thus with n . By Upwards- N -Closure, states that contain a normative converse state of n are normative. Since we consider converse states of n , it is plausible that they are wholly relevant to P . Later we will see that some propositions might have trivial positive or negative content. In *those cases*, one might allow for normative falsemakers without normative truthmakers and vice versa. Yet, since these cases are not regular, we might set them aside for now.
- 23 Note that I do not use the notion of regular propositions in the sense that Fine uses it in [Fine \(2017a\)](#). My way of using it corresponds to what [Fine \(2021\)](#) calls *uniform propositions*.
 - 24 [Fine \(2017a\)](#), p. 461).
 - 25 See [Fine \(2017b\)](#), p. 680).
 - 26 Note that even if the normative disjunct is false the realized content can be normative. For some normative parts of N 's normative truthmakers might obtain, although the truthmaker as a whole does not obtain. For example, if Sarah is permitted but not obliged to help, while permissions are part of obligations.
 - 27 Note that according to the state-based taxonomy, normative negations are indeed normative. However, the label should not be taken to presuppose this result. The label is compatible with the result that normative negations are descriptive.
 - 28 See, for example, [Brown \(2014\)](#), p. 62): 'If a sentence can sincerely be affirmed by one who rejects ethics wholesale, it is not, in a quite natural sense, an ethical sentence'.
 - 29 [Faroldi \(2014\)](#), p. 86).
 - 30 Faroldi introduces this term as a term of art in order to avoid a commitment to a substantive theory of normativity in more fundamental terms, see [Faroldi \(2014\)](#), p. 86, fn.21).
 - 31 See, e.g., [Fine \(2017b\)](#).
 - 32 For a survey, see [Edgington \(2020\)](#).
 - 33 Discussing counterfactuals is beyond the scope of this book. However, truthmaker-based semantics for counterfactuals in [Fine \(2012a\)](#) suggest that the taxonomic results in this section might be extended in a quite natural way.
 - 34 The example is discussed in [Fantl \(2006\)](#), p. 27).
 - 35 An alternative way to interpret these cases is to claim that the normative predicate *good* is predicated to the complex state of affairs in which one's hand is holding a loaded gun aimed at a harmless person and one moves one's trigger finger in a shooting motion, see [Fantl \(2006\)](#), p. 26). Under this reading the proposition plausibly has only normative truthmakers. Consequently, it would be strictly normative.
 - 36 I am following Fine's notation.
 - 37 More recently, [Jago \(2020\)](#), p. 4) develops a semantics for the conditional in relevance logic: $s \Vdash A \supset B$ iff, for all states u such that $u \Vdash A$, $s \sqcup u \Vdash B$. However, the definition is too demanding for our purposes. Take, for example, the conditional 'If Ann promised to ϕ last Monday, then she ought to ϕ '. The definition requires that each exact truthmaker of the antecedent is part of an exact truthmaker of the consequent. Every exact truthmaker of the antecedent must contain a state to the effect that the promise was given *on a Monday*. However, a state that contains such a part plausibly fails to count as an *exact* truthmaker of the proposition [Ann ought to ϕ].
 - 38 Let f be a function mapping each $s \Vdash A$ to a $t_s \Vdash B$. $u \Vdash A \supset B$ iff for some f : $u = \sqcup \{s \rightarrow t_s : s \Vdash A\}$.

- 39 Note that in contrast to $D \rightarrow N$, $D \supset N$ is not entailed by $\neg D$ and thus does not even threaten a simple logical autonomy thesis.
- 40 Note that a state that contains the state that Andrea ought to be reverent does not count as an exact truthmaker of the corresponding green-shoes-conditional.
- 41 See, e.g., [Jago \(2020\)](#).
- 42 See [Fine \(2021, p. 891\)](#).
- 43 On a similar note, see [Schurz \(1997, p. 11\)](#).
- 44 The distinction suggested here is motivated by a distinction between *worldly* and *non-worldly* propositions introduced by [Fine \(2005\)](#). Roughly, a worldly proposition depends upon the circumstances but in a way that it is true *in all* circumstances. A non-worldly proposition does not depend on the circumstances and is true *regardless* of them. An example for a non-worldly proposition is $N \supset N$.
- 45 This is the taxonomic consequence in, e.g., [Oddie \(2018\)](#).
- 46 Another way to go is to adopt a stricter notion of entailment that excludes vacuous entailments. Since this strategy is not based on the classification of propositions, I discuss it in part II of this book.
- 47 Candidates are conceptual, metaphysical, or normative necessities. See [Fine \(2002\)](#) for an elaborated defense of normative necessity. See [Lange \(2018\)](#) for important objections.
- 48 For discussion, see, e.g., [Lange \(2018\)](#). A notable exception are [Cuneo and Shafer-Landau \(2014\)](#) who defend the view that moral fixed points possess conceptual necessity.
- 49 See [van Fraassen \(1969\)](#) for the general idea. Recently, in unpublished works, Johannes Korbmacher und Giuliano Rosella develop truthmaker semantics for modal operators along these lines. See also [Kim \(2024\)](#).
- 50 It should be mentioned that other forms of necessary conditionals such as $\Box((OP \wedge OQ) \supset OP)$ face similar difficulties like the ones discussed in the previous section. Strategies to deal with the non-modal case are arguably applicable to the modal case as well.
- 51 This is a modified version of [Karmo's \(1988\)](#) example.
- 52 Another example that might help to motivate the taxonomic result is the proposition [Every one of the ten commandments is true] which is intuitively normative and structurally similar to P1.
- 53 Note that on fallibilist accounts of knowledge the operator fails to be factive. The assessment of the corresponding proposition then depends on whether those accounts analyze them in a way so that truthmakers of the proposition in the scope of the knowledge operator are exactly relevant to the truth of the knowledge-proposition. If so, the proposition would be normative.
- 54 For example, they give rise to well-known puzzles of ground: they provide *prima facie* counterexamples to the irreflexivity of ground. See, e.g., [Fine \(2010\)](#).

Part II

Logical Autonomy Theses



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

4 Inference Barriers and Modal Gaps

The Conservativeness Explication

A simple explication of the logical autonomy thesis according to which descriptive premises do not entail normative conclusions is subject to several counterexamples. For instance, no matter whether mixed disjunctions are classified as normative or descriptive, they generate counterexamples because D entails $D \vee N$ and $\neg D, D \vee N$ entail N . Let me suppose for the sake of argument that the logical autonomy thesis can be stated in a way in which it is not subject to counterexamples. The logical autonomy debate not only centers around the question of how to explicate a logical autonomy thesis, but is also about its peculiarity. That includes the question in virtue of which feature normativity is logically autonomous. Two extremes span the range of possible responses: on the one hand, the normative domain could be autonomous in virtue of a general feature of logic that is *unrelated* to normativity, on the other hand, it could be autonomous in virtue of characteristic features of normativity that might even be *unique* to the normative domain. Accounts that adopt the first kind of response argue that the normative autonomy thesis is merely one instance of a general principle of logic. The *conservativeness principle* of logic says that you cannot get out of a valid inference what you haven't put in.¹ Accordingly, any two separable domains are autonomous with respect to each other. Applied to the normative domain we have: you cannot validly infer a normative conclusion, if you haven't put in something normative. The normative autonomy thesis thus explicated is not peculiar to normativity at all. Before scrutinizing the non-peculiarity position, it is important to acknowledge that Prior-style examples also refute the conservativeness principle. So, I start by considering two strategies to refine the principle in order to make it immune against Prior-style counterexamples.

Non-Vacuous Entailment

A resolute statement of the non-peculiarity of normative autonomy has been defended in Pigden (1989, 2010a).² Pigden admits that Prior's examples show that the *conservativeness principle* needs to be refined. Still, normative autonomy is just one instance of the refined *general principle*. To see that the structure of Prior's examples can be used to refute the principle that you cannot get out of a valid inference what you haven't put in, consider the following two examples:

- (**Disjunction**) The number 2 is prime. Thus, the number 2 is prime or the moon is made of green cheese.
 (**Conditional**) All tapirs are mammals. Thus, if all mammals love chocolate, tapirs love chocolate.

The examples illustrate that we do not need to put in any green-cheese- or chocolate-information to validly infer green-cheese- or chocolate-conclusions. Many, including Prior himself,³ have recognized that those parts of the conclusion that we haven't put in occupy positions that admit for free variation without affecting the inference's validity. The 'new' parts are *irrelevant* to the inference's validity. To account for this observation, Pigden defines the notion of *inference-relative-vacuity*. His account is based on a notion of classical logical consequence according to which in a valid inference the conclusion is true under all interpretations under which all premises are true.⁴

Vacuous Occurrences: An expression ϕ [predicate or propositional variable] occurs vacuously in the conclusion of a valid inference $K \models X$, iff we can uniformly substitute for any ϕ any other expression ψ of the same grammatical type, yielding a new sentence X' , such that $K \models X'$ is also a valid inference.⁵

Based on this notion of vacuous occurrence in a given inference, Pigden refines the conservativeness principle as follows:

Refined Conservativeness_{VAC}: A predicate or propositional variable cannot occur non-vacuously in the conclusion of a valid inference unless it appears among the premises.⁶

The refined conservativeness principle is immune to the previous examples. Let me illustrate the mechanism by means of the initial normative case: Tea-drinking is common in England \models Tea-drinking is common in England or all New-Zealanders ought to be shot. The normative disjunct

occurs vacuously in the given inference. One could have equally inferred that tea-drinking is common in England or the number 2 is prime. The example does not refute the refined conservativeness principle. In particular, it is compatible with the logical autonomy of the normative domain understood in terms of the refined principle.

It is important to acknowledge that an autonomy thesis that is construed as the refined conservativeness principle is *by design* applicable to all separable domains D_1 and D_2 .⁷ Consequently, for any such domains D_1 and D_2 it holds that D_1 is autonomous with respect to D_2 and vice versa.⁸ The conservativeness explication provides us with an inflationary sense of autonomy. The hedgehog-domain is autonomous from the non-hedgehog-domain in the explicated sense.⁹ That is, we cannot infer hedgehog-conclusions from non-hedgehog-premises unless the hedgehog-expressions occur vacuously in the conclusion of the inference. The same holds for the astronomy-domain, the mammals-domain, etc.

The non-peculiarity of the corresponding normative autonomy thesis seems objectionable. It seems absurd to investigate or scrutinize the autonomy of hedgehogery. By contrast, people take the normative autonomy thesis to be an interesting and important claim about normativity that is worth the debate.¹⁰ However, Pigden argues that the apparent relevance of the normative autonomy thesis is misguided. According to him, people are wrong in thinking that the normative domain is logically autonomous while other domains of discourse are not. He claims that the fact that people have denied normative autonomy explains why there is a debate. According to his reasoning, denying a truism might deserve attention.¹¹ On this view, the significance of the normative autonomy thesis is exhausted by the ignorance of those who deny it.

Next, I discuss another refinement of the initial conservativeness principle that is based on a stricter notion of entailment and truthmaker semantics. I show that a truthmaker-based interpretation of the conservativeness principle is immune to spurious counterexamples. However, being subject to counterexamples is not the only way in which explications of a logical autonomy thesis can be flawed. Even if an explication is not subject to counterexamples, it might fail to specify an interesting sense of autonomy. Eventually, I argue that both conservativeness explications fail in this second regard.

Analytic Entailment

The conservativeness of logic is informally characterized as follows: logic is conservative in the sense that in a valid inference the premises already contain the conclusion. It has been argued that this characterization cannot be taken too literally. Pigden puts it as follows: ‘But venerable and vigorous

as the doctrine is, it remains at the level of *metaphor*. And in logic, at least, we should hope for the literal truth'.¹² Fortunately, the tools of truthmaker semantics provide us with the means to give a valid *literal* interpretation of the informal characterization. We just need to consider a stricter notion of entailment. In the previous chapter, the state-based taxonomy already made use of the notion of conjunctive content parts contained in a proposition P . I propose to interpret the conservativeness principle in terms of *analytic entailment* that is defined in terms of content containment.¹³

Analytic entailment: Propositions P_1, P_2, \dots analytically entail a proposition Q iff $P_1 \wedge P_2 \wedge P_3 \dots$ contains Q .

Recall that P contains Q iff (i) every state that makes Q true is contained in a state that makes P true and (ii) every state that makes P true contains a state that makes Q true. Conjunctions such as $P \wedge Q$ contain each of their conjuncts. Thus, $P \wedge Q$ analytically entails P (Q). In contrast, $P \vee Q$ is usually not contained in P . For $q \Vdash Q$ makes $P \vee Q$ true, but is usually not contained in any of P 's truthmakers. Hence, P does not analytically entail $P \vee Q$. Consequently, a conservativeness principle in terms of analytic entailment is not subject to disjunctive counterexamples. Let me state the refined conservativeness principle and show that the normative domain is autonomous with respect to the descriptive domain thus understood:

Refined Conservativeness_{ANA}: You cannot analytically infer what you haven't put in, i.e., what is not already contained in the premises.

The proof that no descriptive propositions analytically entail a normative proposition is straightforward: let N be a normative proposition. Suppose that P contains N . Since N is normative, it has at least one normative truthmaker. Content containment requires that every state that makes N true is contained in a state that makes P true. Thus, some truthmaker of P contains a normative state. By Upwards- N -Closure, the state in question is normative. Hence, P has a normative truthmaker which means that P is normative. Consequently, every proposition that contains a normative proposition is normative. Note that the proof can be generalized to all domains \mathcal{D} that satisfy Upwards- \mathcal{D} -Closure. Accordingly, the hedgehog-domain is autonomous with respect to its complement domain according to an autonomy thesis in terms of analytic entailment. A conservativeness principle that is based on analytic entailment is immune to counterexamples.

One might ask whether all domains that are autonomous in Pigden's sense are autonomous in the analytic-entailment-sense and vice versa. Pigden's account of autonomy merely requires that a domain can be separated from its

complement on syntactic grounds. Thus, the domain of bachelor-sentences is autonomous with respect to the domain of unmarried-and-male-sentences because sentences in the latter domain do not involve the expression 'bachelor'. Being autonomous in the analytic-entailment-sense is a bit more demanding, for it is based on features of states that make the relevant propositions true. It is not plausible that states that make bachelor-propositions true are in the required sense distinct from states that make unmarried-and-male-propositions true. Thus, the second refinement of the conservativeness principle provides us with a slightly more demanding sense of autonomy. This leads to the question of whether conservativeness explications specify an interesting sense of autonomy.

Autonomy as a Feature of Logic

According to both refinements of the conservativeness principle, the normative domain is logically autonomous with respect to the descriptive domain. However, neither the normative domain nor its relation to the descriptive domain are particularly interesting in virtue of this result. The hedgehog-domain, the domain of astronomy, and in general all suitably separable domains are autonomous with respect to their complement domains.

On these accounts, the autonomy of the normative domain (and any other domain) reduces to the conservativeness of an entailment relation. Both non-vacuous entailment and analytic entailment are conservative. The category of conservative entailment relations allows me to draw an informative distinction between different notions of entailment. Non-vacuous and analytic entailment fall on the one side of the divide; classical entailment falls on the other side. To investigate which entailment relations fall on which side, however, makes any reference to normativity dispensable. To be fair, accounts that answer to the peculiarity question by denying the peculiarity of normative autonomy insist that this is what their considerations show. Logic is concerned with the validity of arguments in virtue of structural features of premises and conclusions irrespective of their concrete contents. Thus, logical autonomy theses must be based on structural features of domains.

The evaluation of an autonomy thesis that employs a conservative entailment relation depends on two aspects. First, one has to assess whether the entailment relation in question is significant in the context of an autonomy debate. Second, since conservativeness is a feature of an entailment relation based on a general feature of domains, namely *separability*, one has to assess whether separability in itself is a significant feature of domains. Both outlined refinements of the conservativeness principle have shortcomings on at least one of these aspects. The second refinement based

on analytic entailment focuses on a quite strict notion of entailment. Note, for example, that given common assumptions about the semantics of existential quantifications, *Fa* does not analytically entail $\exists xFx$.¹⁴ One might hesitate to adopt such a narrow notion of entailment because the significance of a corresponding autonomy thesis is limited. The first refinement is based on separable domains that do not share their characteristic vocabulary. This criterion, however, is easy to meet: the hedgehog domain and the *erinaceidae* domain (the Latin name for hedgehogs) do not share their vocabulary. Again, this criterion is arguably too easy to meet to give rise to a significant sense of autonomy.

I conclude that conservativeness explications specify merely a weak sense of autonomy. Note that this result does not license the conclusion that there is no *significant* sense in which the normative domain is logically autonomous with respect to the descriptive domain. The previous discussion shows that an autonomy thesis should not be based on too strict an entailment relation, if it is supposed to specify an interesting sense of autonomy. Moreover, the distinction between the descriptive and the normative domain should be based on characteristic features of the domains. In the previous chapter, I have argued that the state-based taxonomy provides us with significant distinctions between the two domains. This leaves me with the task to specify an entailment relation so that an autonomy thesis can deal with spurious counterexamples without becoming an overly weak claim. In the next section, I discuss existing logical autonomy theses that illustrate different strategies to deal with counterexamples while specifying stronger senses of autonomy.

The Inference-Barrier Explication

According to one strand in the literature, the logical autonomy thesis is a theorem of the semantics of normative propositions. According to these approaches, autonomy is a feature of normative content. The inference barrier between descriptive and normative propositions is a consequence of a peculiar, though not unique feature of normative content and its truth-conditions. The explications of a logical autonomy thesis discussed in this section are based on taxonomies that I have argued against in [Chapter 2](#). The corresponding inference barriers, however, are important in themselves because they exemplify competing strategies to deal with putative counterexamples. On the one hand, *unrestricted inference barriers* hold without exceptions, but their domain of application is limited; on the other hand, *restricted inference barriers* specify an exception clause, but apply to all propositions. While one could think that the strategies are equivalent, I argue that the accounts differ with respect to their explanatory capacities. I defend the exception-clause-strategy. However, existing accounts based on

this strategy are in danger of trivializing the autonomy thesis. Roughly, the theses easily become sophisticated variants of the trivial statement that no normative proposition is entailed by descriptive propositions unless it is.

Unrestricted Inference Barriers

Russell and Restall (2010) introduce an unrestricted inference barrier. Recall that according to their taxonomy, normative propositions are sensitive to variations of ideal worlds while descriptive propositions are stable under such variations. That is, the choice of ideal worlds is relevant to the truth of normative propositions, yet not to the truth of descriptive propositions.¹⁵

It is crucial for any logical autonomy thesis that it can handle spurious counterexamples like the disjunctive case. Unrestricted inference barriers hold on to the claim that *no* normative proposition is entailed by *any* descriptive propositions. It is the taxonomy that bears the theoretical costs for such an *unrestricted* thesis. It excludes problematic cases from the descriptive and the normative domain. For example, mixed disjunctions are neither classified as descriptive nor as normative. Thus, an autonomy thesis does not apply to them. This is the strategy pursued by Russell and Restall (2010). They present the following explication of a logical autonomy thesis:¹⁶

Unrestricted Inference Barrier: No normative proposition is entailed by descriptive propositions, i.e., no proposition that is fragile under normative changes is entailed by propositions that are stable under normative changes.

One advantage of this proposal is that one can *prove* the autonomy thesis. Consider a putative counterexample. Choose an arbitrary model that satisfies the descriptive premises. If the model does not satisfy the normative conclusion, the argument is not valid and one is done. If the model satisfies the normative conclusion, the model can be changed in a way that it does not satisfy the normative conclusion. This is possible because the normative conclusion is fragile under normative changes. Since these changes make no difference to the truth of the descriptive premises, the new model satisfies the premises, but not the conclusion. Thus, the argument is not valid.¹⁷

Another advantage of the proposal is that it provides a unified account of inference barriers that have been discussed in the literature.¹⁸ For example, an inference barrier between *particular propositions* like [Sarah is mortal] and *general propositions* like [All humans are mortal] can be captured by the same general pattern: propositions that are stable under

changes concerning the domain of individuals do not entail propositions that are fragile under those changes.

That said, let me turn to the down-sides. We already saw that the taxonomy is objectionable. Since the autonomy thesis is essentially based on the classifications, it inherits this criticism. Moreover, the explication of the logical autonomy thesis is entailed by semantic and taxonomic assumptions. Thus, views that deny autonomy cannot accept these assumptions. Let me now focus on a different aspect: I argue that unrestricted inference barriers fail to equip us with *explanations* for spurious counterexamples. To see the significance of this point, it helps to make the dialectical situation explicit. Autonomists claim that normativity is logically autonomous. Non-autonomists came up with counterexamples. Setting aside those examples—given that we grant their validity—is only legitimate, if there is *a reason* to think that they do not refute the autonomy of normativity.

For concreteness, let me focus on the disjunctive case $D \neq D \vee N$. The disjunctive case requires an explanation why the fact that D entails $D \vee N$ does not reveal that normativity is not autonomous. Note that unrestricted autonomy theses narrow down the normative and the descriptive domain so that $D \vee N$ belongs to neither category.¹⁹ On the face of it, the strategy suggests that there are no problematic cases. The disjunction $D \vee N$ does not pose a greater threat to a normative autonomy thesis than the fact that [Sarah is Ukrainian] entails [Someone is Ukrainian]. For this strategy to succeed, one must be convinced that the normative autonomy debate is essentially concerned with stable and fragile propositions. Even if we grant that much, the reason why $D \vee N$ is not stable is that it is sometimes sensitive to normative changes. Thus, the disjunction has features of normative propositions. On the other hand, it is entailed by a stable proposition and thus is not fragile. Yet, these are exactly the features that make the example a putative counterexample in the first place: the conclusion is entailed by a descriptive proposition and has some normative features. Now a promising explanation why the example does not reveal that normativity is not autonomous could point out that by the lights of the account, the disjunction has no normative features whenever the conclusion is true in virtue of the descriptive premise.²⁰ However, the unrestricted autonomy thesis does not make use of its explanatory capacities. It simply excludes mixed disjunctions from the normative and descriptive category. By contrast, restricted inference barriers specify exception clauses. I shall argue that this strategy increases their explanatory power.

Restricted Inference Barriers

Restricted inference barriers are not based on limited categories, but on exhaustive categories. Consequently, an inference barrier does not hold

without exceptions and thus requires an additional clause specifying those exceptions. The key feature of restricted inference barriers is the plausibility of their exception clause. To see why, note that there are trivial ways to immunize an autonomy thesis against counterexamples. For instance, pointing out that no normative proposition is entailed by descriptive propositions *unless* it is entailed by descriptive propositions is trivial. To avoid trivialization and question begging, it is crucial that an exception clause provides an explanation why these cases should be excluded. To get a better grip on the strategy and impending threats, let me introduce two restricted inference barriers and compare the proposals. The first variant is developed in [Singer \(2015\)](#), the second variant is developed in [Russell \(2021\)](#). Recall that Singer defines descriptive propositions in terms of norm-invariance and normative propositions in terms of norm-variance. An *S*-shift on Russell's account requires to change the set of ideal worlds in a model.

Restricted Inference Barrier_{NORM}: Satisfiable descriptive premises Γ do not entail a normative conclusion C , unless $\Lambda\Gamma \wedge C$ is norm-invariant.²¹

Restricted Inference Barrier_{MODEL}: Non-normative Γ do not entail normative C , unless all *S*-shifts of all models of Γ are also models of C .²²

Let me first explain how the exception clauses are motivated. Singer argues that some inferences from descriptive premises to a normative conclusion are not normatively relevant. For an inference to be normatively relevant the conclusion's normative impact must be relevant to the possibilities reasoned about.²³ Let me illustrate the proposal by going through the disjunctive case. If we reason about *D*-possibilities, the normative impact of $D \vee N$ is—so the argument goes—not relevant because $D \vee N$ is true anyway. Note that on Singer's account, normative impact is measured in terms of norm-variance. Thus, according to his account, one has to consider the conjunction of the premises and the conclusion and see whether the conjunction is norm-variant, i.e., has normative impact. If the conclusion does not have normative impact with respect to the possibilities reasoned about, we can legitimately exclude the inference for being normatively irrelevant according to Singer. Russell, on the other hand, argues that there are well-known counterexamples to a class of inference barriers. All of them can be explained by the same mechanism. All possible *S*-shifts of models of the premises are models of the conclusion. The motivation here seems to be the unifying capacity of the exception clause.

I shall now turn to the proofs of the theses. The proofs will help me to assess whether the exception clauses are justified, for they reveal how

the exception clauses work. Let me start with Singer's account: Singer assumes an intensional account of propositions. On an intensional account, the conjunction of premises and the conclusion that is entailed by those premises *just is* the conjunction of the premises. The reason is that the conjunction of conclusion and premises is necessarily equivalent with the conjunction of the premises. Consequently, whenever norm-invariant premises entail any conclusion, the conjunction of the premises and the conclusion is by assumption norm-invariant. Thus, any example where descriptive premises entail a normative conclusion meets the exception clause.

Similar reasoning applies to Russell's proposal. If non-normative Γ entails normative C , then all of Γ 's models are models of C . Since Γ is stable under S -shifts all S -shifts of models of Γ are models of Γ . Thus, they are models of C which means that the exception clause is satisfied. Again, any example where descriptive premises entail a normative conclusion meets the exception clause.

Despite their structural similarities, the two accounts are problematic for different reasons. The second barrier thesis is almost the trivial exception clause I began with. To see why, note that being an S -shift of a model is a reflexive relation. In other words, every model is an S -shift of itself. Thus, if every S -shift of a model of Γ is also a model of C (exception clause), then every model of Γ is a model of C . Hence, Γ entails C . Together with the previous proof, this result entails that the exception clause is equivalent to the proposition that descriptive Γ entails C . The justification for the exception clause in the second proposal states that the normative impact of a conclusion must be relevant to the possibilities reasoned about. However, the result that $\Lambda\Gamma \wedge C$ has no normative impact in worlds where $\Lambda\Gamma$ is true is due to the fact that intensional accounts are too coarse-grained to distinguish between the normative impact of $\Lambda\Gamma$ and the normative impact of $\Lambda\Gamma \wedge C$. Any reasonable account of normative impact should account for the normative impact of normative propositions irrespective of whether they are entailed by descriptive premises or not.

I conclude that restricted inference barriers are advantageous over unrestricted inference barriers because exception clauses can explain why spurious counterexamples arise in the first place and why they are indeed spurious. The clauses explain why unrestricted barriers fail when they are applied to the whole normative domain. However, existing accounts do not provide satisfying explanations why the excluded cases are spurious rather than genuine counterexamples. Clauses that do not go beyond the claim that a descriptive proposition entails a normative proposition do not provide explanations at all.

Autonomy as a Feature of Semantics

Let me return to the peculiarity question I began with. The discussed accounts assume that there is a component in the semantic framework that is relevant to the truth-evaluation of normative, yet not descriptive propositions. The autonomy of normativity is in turn defined in terms of modal independence of the semantic components. These explications of the logical autonomy thesis are neither universal nor unique. They are not universal because not every arbitrary domain requires an independent component in the semantic framework. For example, hedgehog- and erinaceidae-propositions are arguably evaluated with respect to the same components in the semantic framework. The explications are not unique because other domains arguably share this feature. For example, general propositions depend on a domain of individuals in a way that particular proposition do not. Thus, there is an inference barrier between the two domains. This is a satisfying result. That said, it is important to bear in mind that the inference barriers depend on the fact that the truth-conditions of normative propositions require a modally independent semantic component. This assumption is controversial. It is in danger of replacing the autonomy debate. One thus faces a debate about the correct semantics instead of an autonomy debate.

This result is in tension with the authors claim that their account works with *standard* semantics, for standard semantics should be available for autonomists and non-autonomists alike.²⁴ An impartial semantic framework and taxonomy, however, are compatible with putative counterexamples. A satisfying explication of the logical autonomy thesis thus depends on a plausible explanation why particular counterexamples are spurious rather than genuine. I have argued that restricted inference barriers have the means to provide such explanations in terms of their exception clause. A compelling exception clause must be explanatory. In the next section, I discuss a state-based explication of the logical autonomy thesis developed in [Fine \(2021\)](#). The account shows that we can explain spurious counterexamples without making commitments to modally independent semantic components.

A Gap Through Fine-Grained Attribution

[Fine \(2021\)](#) develops a state-based explication of the logical autonomy thesis which he calls the *Gap Principle*. The exception clause of the Gap Principle is based on *attributions* of modal relations between descriptive and normative states, in particular, incompatibilities between them and their descriptive parts. This way of understanding the autonomy thesis is in many respects illuminating. One significant result of Fine's discussion

is a nuanced distinction between the notions of *autonomy*, *independence*, and *separation* of the normative domain. The fine-grained distinctions are indispensable for a modal explication of the autonomy thesis.²⁵ However, I shall argue that the proposal faces a serious objection. Views that deny the autonomy of the normative domain can accept the Gap Principle provided that they accept intimate connections between descriptive and normative states. This, however, raises doubts concerning the joint-carvingness of the Gap Principle.

The Gap Principle

Let me remind the reader of distinctions established by the state-based taxonomy. A proposition is normative iff it has at least one normative truthmaker. A proposition is descriptive iff it has only descriptive truthmakers. Before introducing the Gap Principle, I shall introduce the relevant concepts and notational conventions. First, we need the concept of classical entailment. Recall that within possible world semantics, P entails Q iff $P \subseteq Q$. The analogue does not work on an account of exact truthmaking. For example, $P \wedge Q$ entails Q , but $(P \wedge Q) \not\subseteq Q$ because $p \sqcup q \Vdash P \wedge Q$ does not *exactly* make Q true. Instead, we shall say that P entails Q iff every state that is incompatible with Q is incompatible with P , where a state s is incompatible with a proposition P iff it is incompatible with all of its truthmakers. This notion corresponds to the classical notion of entailment. Moreover, for an arbitrary proposition P , $\sim P$ is the set of all consistent states incompatible with P .²⁶ Furthermore, let P_D denote P 's restriction to its consistent descriptive content. For each of P 's truthmakers we take its consistent maximal descriptive part.²⁷ P_D is then the set of these consistent descriptive states. For example, let P be the proposition that Sarah ought to help or it is rainy. P_D comprises P 's descriptive truthmakers such as that it is rainy and the descriptive parts of P 's normative truthmakers such as that Sarah exists. Fine proposes the following explication of the logical autonomy thesis:

Gap Principle: No descriptive proposition P entails a normative proposition Q unless $(\sim Q)_D$ entails $\sim P$, i.e., unless the descriptive component of every possible state incompatible with Q is incompatible with P .²⁸

Note that the Gap Principle is a restricted inference barrier. To provide the reader with a better grasp of the Gap Principle, I will go through three putative counterexamples to an autonomy thesis discussed by Fine. The examples illustrate the explanatory power of the proposal. On the face of it, the examples seem to be based on different mechanisms, but the Finean

analysis shows that they are in fact based on the same mechanism. Moreover, the examples illustrate that the exception clause provides an informative explanation of what it is that makes counterexamples spurious.

- (Disjunction) Tea-drinking is common in England. Thus, tea-drinking is common in England or Mary ought to attend the meeting.
- (Non-Existence) Mary does not exist. Thus, it is not the case that Mary ought to attend the meeting.
- (Obligation-Permission) Mary exists. Thus, Mary ought to attend the meeting or it is permitted that Mary does not attend the meeting.

First, consider (Disjunction). Any state *s* that is incompatible with the conclusion must be incompatible with both disjuncts. Thus, *s* is incompatible with the premise because the premise is one of the disjuncts. Accordingly, the premise entails the conclusion. The reason why *s* is incompatible with the premise [Tea-drinking is common in England] is that it contains descriptive parts, e.g., that tea-drinking is rare in England, that rule out the descriptive disjunct. Its normative parts are redundant with respect to the incompatibility with the premise. This is what the exception clause tells us. If we restrict our attention to the consistent descriptive parts of states that are incompatible with the conclusion and discover that these descriptive parts are still incompatible with the premises, then the entailment relation holds *in virtue of* the descriptive parts. To put it in other words, the incompatibility is *attributable* to the descriptive components.

Let me turn to (Non-Existence). States that are incompatible with the conclusion contain states to the effect that Mary ought to attend the meeting. The state that Mary ought to attend the meeting contains the descriptive state that Mary exists as integral part. Again, the fact that states that are incompatible with the conclusion are incompatible with the premise is due to a descriptive part of those states, namely the state that Mary exists. The incompatibility is attributable to a descriptive component.

Finally, in (Obligation-Permission), the sole consistent state that is incompatible with the conclusion is the descriptive state that Mary does not exist. Thus, the incompatibility with the premise is attributable to a descriptive state. This shows that all three examples work by the same mechanism. The fact that the premises are incompatible with states that would rule out the conclusion can be traced back to their descriptive components.

The account suggests the following characterization of spurious counterexamples: the entailment relation between premises and conclusion ‘can be seen to hold for reasons having nothing to do with the normative content’ of the conclusion.²⁹ And on the same page, concerning genuine

counterexamples: ‘we did not want there to be a descriptive route from *P* to *Q*’. Fine claims that by excluding all cases where the entailment relation is attributable to the descriptive components, we specify a sense in which the normative domain is logically autonomous with respect to the descriptive domain. Fine emphasizes that the Gap Principle highlights that autonomy thus understood is less demanding than many extant accounts presupposed. To see why, I shall introduce Fine’s distinction between autonomy and two other principles on the level of states:³⁰

- Autonomy:** If a descriptive state *s* is incompatible with a possible non-descriptive state *t* then *s* is incompatible with a descriptive part of *t*, i.e., the incompatibility is attributable to some descriptive part of *t*.³¹
- Independence:** Any descriptive state is compatible with any purely normative state.
- Separation:** Any state is either descriptive or is the fusion of a descriptive and a purely normative state.

The discussion of existing taxonomies and logical autonomy theses revealed that many accounts are based on the assumption that, for example, the semantic framework contains a normative component that is modally independent from descriptive components. Moreover, this assumption requires that the normative and the descriptive components can be *separated*. I have argued that these assumptions face serious objections in the context of an autonomy debate. Similarly, the principles *Independence* and *Separability* require the existence of purely normative states. However, the existence of purely normative states can reasonably be questioned. Even atomic normative proposition such as [Sarah ought to help] or [Theft is morally wrong] seem to involve descriptive components. It is a great advantage of the Finean account that it reveals that neither *Autonomy* nor the Gap Principle require the existence of purely normative states. That said, I shall now turn to a problem that accompanies all exception clauses which I will call *overgeneralization problem*. To put it in Fine’s own terms: ‘For in throwing out the cheats, we may throw out genuine cases to which the principle is meant to apply’.³²

The Overgeneralization Problem

The Gap Principle *explains* the unifying mechanism behind spurious counterexamples. One difficulty, however, is that the attribution-strategy is in danger of overgeneralizing. I shall argue that views that deny autonomy might be happy to attribute all kinds of incompatibilities between descriptive and normative states to descriptive components. If this is correct, the Gap Principle is acceptable for views that deny autonomy. This suggests

that the Gap Principle does not draw a satisfying line between autonomists on the one hand and non-autonomists on the other hand.

To see why the attribution-strategy is in danger of overgeneralizing, it will be illuminating to first discuss an example that can be construed as a putative genuine counterexample. To start with, I shall consider the example in light of a stronger variant of *Autonomy*. Thereafter, I explain why Fine takes *Autonomy* (as stated above) to respond correctly to the example. Finally, I raise doubts concerning this result. If my reasoning is correct, *Autonomy* can be seen to hold, even though intuitively genuine counterexamples succeed. Let me first introduce the example:³³

(**Promise**) Sarah promised to pay five dollars. Thus, Sarah ought to pay five dollars.

A first thing to note is that a stronger version of *Autonomy* referring to possible worlds is compatible with (Promise) and similar examples and thus overgeneralizes:

World-Autonomy: If a descriptive state s is incompatible with a possible world w then s is incompatible with a descriptive part of w , i.e., the incompatibility is attributable to some descriptive part of w .³⁴

To see that World-Autonomy is compatible with putative genuine counterexamples, note that if we accept (Promise), worlds in which Sarah is permitted not to pay five dollars are worlds in which she did not promise to pay five dollars. Accordingly, these worlds plausibly contain parts to the effect that Sarah did not promise to pay five dollars. Consequently, the descriptive state that Sarah promised to pay five dollars will be incompatible with worlds in which Sarah is permitted not to pay five dollars, but the incompatibility is attributable to descriptive parts of those worlds. Thus, World-Autonomy is compatible with (Promise).

It is important to acknowledge that Fine argues that the weaker principle *Autonomy* is not subject to the same overgeneralization problem. The reason is that *Autonomy* considers states instead of whole possible worlds. According to the argument, the normative state that Sarah is permitted not to pay does not contain a descriptive part that is incompatible with the descriptive state that Sarah promised to pay. Accordingly, *Autonomy* would be refuted by examples like (Promise), which is the desired result because they are putative *genuine* counterexamples.

Let me scrutinize this result. Suppose we accept (Promise). Consequently, we want to deny *Autonomy*. Therefore, we are committed to the claim that the descriptive state that Sarah promised to pay five dollars is incompatible

with the normative state that Sarah is permitted not to pay five dollars. Suppose someone asks us why the states are incompatible. What we cannot do is to point to descriptive parts of the normative state. We cannot say that permissions are partly constituted by the absence of promises to do otherwise because this would suggest that the normative state has some descriptive parts incompatible with the state that Sarah promised to pay five dollars. Attempts to partially analyze the normative content of the conclusion in descriptive terms renders *Autonomy* and (Promise) compatible. To approach it from a different angle, we can describe (Promise) on the level of states by saying that a descriptive state necessitates a normative state. Given that we cannot point to any descriptive components of the normative state, we face a necessary connection between wholly distinct states which might raise the objection that the connection is brute. After all, the descriptive and the normative state are otherwise unrelated. The commitment to brute necessary connections, however, is commonly taken to count against a view.³⁵ To be non-autonomist in Fine's sense, we have to accept brute necessary connections between descriptive and normative states, but it is implausible that we cannot be non-autonomists without accepting such brute necessary connections.

This result is particularly pressing in light of the assumption of non-trivial integral descriptive parts. To see that, it will be illuminating to modify (Promise) so that the conclusion is logically equivalent to the original conclusion: it is not the case that Sarah is permitted not to pay five dollars. Now, recall the example (Non-Existence): Sarah does not exist. Thus, Sarah is not obliged to attend the meeting. Fine claims that Sarah's non-existence can very well be seen as an immediate falsemaker of the proposition [Sarah ought to attend the meeting]. In addition, he claims that the following case can be handled in the same way: [Sarah cannot attend the meeting] entails [Sarah is not obliged to attend the meeting]. Again, the state that Sarah cannot attend is taken to be an immediate falsemaker of the proposition [Sarah ought to attend the meeting]. If we ask why, a reasonable response seems to be that existence and ability are taken to be integral parts of obligations. Now, compare the modified (Promise) case: [Sarah promised to pay] entails [It is not the case that Sarah is permitted not to pay]. Someone who accepts (Promise) might insist that Sarah's promise is an immediate falsemaker of the proposition [Sarah is permitted not to pay]. By some non-autonomist, the absence of promises to ϕ might even be seen as integral part of permissions not to ϕ . If there is a difference in the sequence of examples, it seems at best a matter of degree. To be clear, I do not aim to defend such a view. However, it strikes me as implausible that *non*-autonomists are committed to a strict separation of descriptive and normative states just to be able to *deny* Autonomy.

The main problem becomes transparent when we acknowledge a small difference between two ways in which Fine informally characterized

spurious counterexamples: on the one hand, we want to exclude cases that hold for reasons having nothing to do with normative content, on the other hand, we want to exclude cases with a descriptive route from D to N . These characterizations, however, can come apart. One kind of example that illustrates this claim is the previously discussed case where we accept (Promise), but take descriptive routes to matter to normative content. Another kind of example that illustrates this claim results from structural cases of overdetermination. To see that, let D be a descriptive proposition that entails a normative proposition N and a descriptive proposition D' . Hence, D entails $D' \vee N$. Every state that is incompatible with the disjunction contains a descriptive component that is incompatible with D' and thus with D . The example is compatible with the Gap Principle. The example meets the exception clause because there is a descriptive route from D to $D' \vee N$, though the example also holds for reasons having to do with normative content. So, the two characterizations can come apart. Notice that *Autonomy* would still be false because of D and N . A fully satisfying account of logical autonomy, however, should be able to classify the disjunctive case as genuine counterexample.

I conclude that the exception clause in the Gap Principle overgeneralizes. Next, my aim is to develop an improved proposal by focusing on the first of the outlined informal characterizations. Spurious counterexamples hold for reasons having nothing to do with normative content. That is, the normative content is *irrelevant* to the entailment relation.

Summary and Prospects

In this chapter, I have been concerned with a series of different explications of logical autonomy theses. The simple logical autonomy thesis according to which descriptive premises do not entail a normative conclusion can be developed in different directions. I discussed:

Conservativeness Explications:	The entailment relation in question is conservative, i.e., premises do not entail conclusions that (non-vacuously) involve parts (structural or content parts) not already contained in the premises.
Unrestricted Barrier Explications:	Descriptive premises do not entail normative conclusions. The underlying taxonomy provides narrow categories of descriptive and normative propositions.
Restricted Barrier Explications:	Descriptive premises do not entail normative conclusions unless they meet a particular exception clause.

The discussion revealed that conservativeness explications that are based on a general feature of domains such as separability merely entail a weak sense of autonomy. In this weak sense, nearly every domain is autonomous from several other domains. Thus, these kinds of autonomy theses do not reveal any insights into a particular domain. Moreover, I have argued that unrestricted barrier explications set aside difficult cases. Since the difficult cases are no longer in the scope of an autonomy thesis, the accounts do not provide us with explanations why spurious counterexamples arise and why they are indeed spurious. Restricted barrier explications respond to this challenge by specifying features of spurious counterexamples that must be excluded. However, this strategy faces difficulties. The exception clause is in danger of trivializing the autonomy thesis. In some cases, the formulation of the exception clause turns out to nearly approximate the negation of the unrestricted part of the autonomy thesis. In other cases, the exception clause excludes too much so that the corresponding logical autonomy thesis is acceptable for views that deny the autonomy of the normativity. These considerations motivate two desiderata for a satisfying logical autonomy thesis:

- D1: An autonomy thesis explicates an interesting sense of autonomy that reveals something distinctive about normativity.
- D2: An autonomy thesis draws a line between spurious and genuine counterexamples and explains why the former do not pose a threat to normative autonomy.

In the remainder of the second part of this book, I argue in favor of a particular way to respond to the desiderata. I have argued that the Finean Gap Principle provides us with a strong explanation for the mechanism that unifies spurious counterexamples. I have also argued that the focus on descriptive routes might be too strong because we thereby exclude cases that seem to be putative genuine counterexamples. However, the motivating idea that spurious counterexamples hold even though the conclusion's normative content is irrelevant for the entailment relation is promising. Thus, a logical autonomy thesis must be concerned with *relevant entailments*. The discussion so far revealed that even weak senses of autonomy require relevance constraints. I shall discuss whether a logical autonomy thesis might eventually require relevance constraints on two dimensions. [Chapter 5](#) is concerned with an account of partially irrelevant conclusions. I develop a state-based account of relevant entailments and propose a modified Gap Principle, namely the *Re-Gap Principle*. I argue that the Re-Gap Principle provides us with a satisfying account of the modal autonomy of normativity. In [Chapter 6](#), I turn to another dimension of normative relevance. Against the background of putative exceptions to the

Re-Gap Principle, I consider a restriction of the principle to authoritative normativity. The restricted Re-Gap Principle is a way to maintain a substantive sense in which the normative domain is modally autonomous in case the considered exceptions were to succeed.

Notes

- 1 Note that the notion of conservativeness often shows up in logic, when talking about conservative extensions of a theory. Those extended theories do not prove new theorems about the language of the original theory. This technical sense of conservativeness is not the sense in which I am using it here.
- 2 For an early statement of the view that normative conclusions cannot be derived from propositions that do not contain normative concepts, see [Bolzano \(1837, §200\)](#).
- 3 [Prior \(1960, p. 203\)](#).
- 4 The account applies to classical propositional and first-order predicate logic with primitive predicates. Pigden's proposal is committed to the denial of the claim that 'ought' or 'permission' are sentential operators. Otherwise, there might be deontic-logical truths such as $P(A) \vee P(\neg A)$ that are entailed by arbitrary descriptive sentences, while replacing the sentential operator 'P' affects the validity of the inference, see [Pigden \(1989, p. 136\)](#).
- 5 [Pigden \(1989, p. 136\)](#).
- 6 [Pigden \(1989, p. 136\)](#).
- 7 For our purpose, it is not important that we have a formally precise notion of separability. Roughly, two separable domains are non-overlapping sets of sentences where one of them contains all sentences that involve particular non-logical expressions (of a background language), while the other domain contains only sentences that do not involve those expressions.
- 8 For a critical assessment of this consequence, see, e.g., [Maguire and Woods \(2017\)](#).
- 9 See [Pigden \(1989, §9\)](#).
- 10 See, for example, the accounts discussed in this book. See also [Humberstone \(2019\)](#), for a rich and helpful discussion of several interpretations of the logical autonomy thesis.
- 11 See [Pigden \(1989, pp. 146–147\)](#).
- 12 [Pigden \(1989, p. 130, my italics\)](#).
- 13 See [Fine \(2016\)](#) for elaborated work on a logic of analytic entailment that goes back to [Angell \(1989\)](#). Fine shows that an autonomy thesis in terms of content containment holds, see [Fine \(2021, p. 895\)](#). It is noteworthy, however, that we thereby obtain a literal interpretation of the conservativeness principle.
- 14 Note also that this result might bear on the question of whether normativity is autonomous: That Sarah exists might entail that there is something that is obliged to ϕ or is permitted not to ϕ .
- 15 This is a slightly simplified version of their definition, see [Russell and Restall \(2010, pp. 254–255\)](#) and sec. 'Changing Ideal Worlds' in this book.
- 16 For the formal statement and a detailed proof, see [Russell and Restall \(2010, pp. 255–256\)](#).
- 17 Notice that Russell and Restall's taxonomy applies only to satisfiable propositions.
- 18 See [Russell and Restall \(2010, p. 248\)](#).

- 19 [Wolf \(2020, p. 89\)](#) objects that this strategy to deal with counterexamples just gives up the game.
- 20 Note that I do not take the explanation for disjunctive cases to be fully satisfying. In cases where D and N are true, $D \vee N$ seems to have normative features, but the account cannot account for this assessment. Later I will turn to better explanations.
- 21 See [Singer \(2015, p. 203\)](#). Singer does not state his explication by means of an exception clause, but his formulation is equivalent to the one stated here. His informal motivation suggests that he wants to exclude the normatively irrelevant cases.
- 22 See [Russell \(2021, pp. 23–24\)](#).
- 23 [Singer \(2015, p. 203\)](#).
- 24 That the accounts claim to work with standard semantics for normative expressions has been emphasized and discussed in [Chapter 2](#).
- 25 In [Chapter 5](#), I develop a principle, the Re-Gap Principle, that is based on the Finean insights. I shall discuss how the Finean Gap Principle relates to the Re-Gap Principle in that chapter.
- 26 In previous chapters, I used ‘ \sim ’ to refer to internal negation. Since we are not concerned with internal negations henceforth, in the remainder of this book ‘ $\sim P$ ’ should be read as it is defined here.
- 27 d is the maximal descriptive part of s iff (i) d is descriptive, (ii) $d \sqsubseteq s$, and (iii) d contains all states that satisfy (i) and (ii). See [Fine \(2021, p. 898\)](#).
- 28 [Fine \(2021, p. 899\)](#).
- 29 [Fine \(2021, p. 985\)](#).
- 30 [Fine \(2021, pp. 900–901\)](#).
- 31 *Autonomy* and the Gap Principle are equivalent given the assumption that any consistent state is compatible with some normative state, see [Fine \(2021, p. 910\)](#).
- 32 [Fine \(2021, p. 895\)](#).
- 33 This is a modified version of [Searle’s \(1964\)](#) example. It has also been denied that the counterexample is genuine, see [Hare \(1964\)](#). The basis for that reasoning, however, employs resources that importantly differ from the reasoning pointed at spurious counterexamples. I will return to the example in [Chapter 6](#).
- 34 [Fine \(2021, p. 900\)](#).
- 35 In particular, it has been objected to non-naturalists that they are committed to brute necessary connections, see [McPherson \(2012\)](#).

5 The Relevant Gap Principle

Sources of Partial Irrelevance

Let me begin with a brief survey to see where we stand. In the previous chapter, I have argued that the Finean approach provides a convincing explanation why some counterexamples to the logical autonomy of the normative domain are merely spurious. We can legitimately treat counterexamples as merely spurious, if the entailment relation holds for reasons having nothing to do with the conclusion's normative content. As I understand this explanation, it is best construed as the claim that if the conclusion's normative content is *irrelevant* to the entailment relation the counterexample is spurious. However, I have criticized the Finean approach for classifying some putative counterexamples as spurious that are intuitively classified as genuine counterexamples. A plausible diagnosis for the misclassification is that these examples hold in virtue of the conclusion's descriptive *and* its normative content. Thus, the normative content seems to be relevant. On the Finean account, however, the conclusion's normative content is irrelevant in those cases. I have argued that this result is unsatisfying.

These observations suggest that counterexamples fall along a spectrum: some examples are such that the conclusion's normative content is clearly irrelevant to the entailment relation, while in other examples there is no doubt that the conclusion's normative content is highly relevant to the entailment relation. To illustrate both extremes, let me give one example for each. Suppose Sarah promises to donate one percent of her income. Suppose further that as a matter of conceptual truth, keeping one's promises is morally obligatory. Thus, that Sarah promises to donate one percent of her income conceptually entails that Sarah ought to donate one percent of her income. In that case, it is uncontroversial that the conclusion's normative content is relevant to the entailment relation. The entailment holds—at least this is what I assume for the purposes of illustration—in virtue of a conceptual connection between normative and descriptive contents. By

contrast, consider the proposition that Macron is the French president. It entails that Sarah ought to donate or Macron is the French president. It is evident that in this case the conclusion's normative content plays no role and is thus irrelevant to the entailment relation. On the account that I am going to defend, this explains why the latter but not the former example is a spurious counterexample.

A central aim of this chapter is to draw an accurate line between spurious and genuine counterexamples with respect to cases that fall in-between the extremes. That is, cases where it is difficult to tell whether the conclusion's normative content is irrelevant to the entailment relation. To that end, we need to give an account of a fine-grained notion of *irrelevance* that can accommodate for that difficulty. Before turning to that aim, I shall substantiate the intuitive assessment that examples falling on the irrelevance-side of the spectrum do not show anything distinctive about normativity. Thereafter, in a brief aside, I show that irrelevant conclusions generate trouble in several debates. The digression helps to overcome concerns according to which spurious counterexamples suggest that there is no substantial sense in which normativity is autonomous.

A Cartography of Possibilities

It is a natural thought that we can settle the modal autonomy question (i.e., whether the normative domain is modally autonomous with respect to the descriptive domain) by settling the totality of modal relations between the normative and the descriptive. Once the modal profiles of the domains are fixed, they indicate whether any descriptive premises entail a normative conclusion. In this section, I am going to show that some spurious counterexamples to modal autonomy theses are valid, even though the modal profile of the normative can be reasonably described as *independent* from the descriptive. This result substantiates the intuitive assessment that some spurious counterexamples do not tell us anything about normativity and its relation to the descriptive. As will become clear, their defectiveness in this respect is due the conclusion's normative content being irrelevant.

Extremes on the irrelevance spectrum can be illustrated by means of a *toy state space*. I shall consider a minimal state space that comprises only four basic states.¹ Admittedly, this is an artificial set-up. However, the considerations to follow illustrate something genuine: some descriptive premises entail a normative conclusion irrespective of the modal profile of normativity. To see that, let me consider the toy state space. Suppose there are two basic normative states that cannot obtain together: $\mathcal{S}^N = \{n, n'\}$. To pick up on the previous example, let n be a state to the effect that Sarah ought to donate and n' a state to the effect that Sarah may keep her money.

Basic normative states are such that we can obtain all normative states in the state space by fusing basic normative states in \mathcal{S}^N or by fusing these normative states with descriptive states. We assume that basic normative states do not contain non-trivial descriptive parts.² Additionally, we have two basic descriptive states that cannot obtain together: $\mathcal{S}^D = \{d, d'\}$. Let d be a state to the effect that Sarah promises to donate and d' a state to the effect that Sarah does not promise to donate.

In a modalized state space, the undertaking of settling the modal profile of the normative and the descriptive can be carried out by settling which states belong to the set of consistent states \mathcal{S}^\diamond . Two states are incompatible with each other iff their fusion is inconsistent, and compatible otherwise. A state s necessitates another state t iff s is incompatible with all states that are incompatible with t . In our toy state space, there are merely four options concerning the modal profile of a basic normative state: (i) the basic normative state is *independent* from the descriptive states, (ii) a descriptive state is *necessary* for the normative state, (iii) a descriptive state is *sufficient* for the normative state, or (iv) a descriptive and the normative state necessarily *co-obtain*. So, we can describe *four* modalized state spaces by modifying \mathcal{S}^\diamond accordingly:

- Independence:** $(\mathcal{S}, \mathcal{S}_I^\diamond, \sqsubseteq)$ where $\mathcal{S}_I^\diamond = \{d \sqcup n, d \sqcup n', d' \sqcup n, d' \sqcup n'\}$. Since all combinations of basic normative and descriptive states are consistent, they are independent. Neither descriptive states fix normative states nor vice versa.
- Necessity:** $(\mathcal{S}, \mathcal{S}_{nC}^\diamond, \sqsubseteq)$ where $\mathcal{S}_{nC}^\diamond = \{d \sqcup n, d \sqcup n', d' \sqcup n'\}$. Since $d' \sqcup n \notin \mathcal{S}_{nC}^\diamond$, n is incompatible with all states incompatible with d . Hence, n necessitates d which in turn means that d is necessary for n to obtain.
- Sufficiency:** $(\mathcal{S}, \mathcal{S}_{sC}^\diamond, \sqsubseteq)$ where $\mathcal{S}_{sC}^\diamond = \{d \sqcup n, d' \sqcup n, d' \sqcup n'\}$. Since $d \sqcup n' \notin \mathcal{S}_{sC}^\diamond$, d is incompatible with all states incompatible with n . Hence, d necessitates n which in turn means that d is sufficient for n to obtain.
- Equivalence:** $(\mathcal{S}, \mathcal{S}_E^\diamond, \sqsubseteq)$ where $\mathcal{S}_E^\diamond = \{d \sqcup n, d' \sqcup n'\}$. Since \mathcal{S}_E^\diamond contains neither $d \sqcup n'$ nor $d' \sqcup n$, d and n necessarily co-obtain.

It is worth acknowledging that we thereby obtain a partial order on the set of \mathcal{S}^\diamond -variants with respect to the subset relation: $\mathcal{S}_E^\diamond \subset \mathcal{S}_{sC}^\diamond \subset \mathcal{S}_I^\diamond$ and $\mathcal{S}_E^\diamond \subset \mathcal{S}_{nC}^\diamond \subset \mathcal{S}_I^\diamond$. This allows one to measure relative degrees of entanglement between the domains. Assuming modal independence of the normative corresponds to a minimal degree of entanglement, while equivalence corresponds to the maximal degree of entanglement.

This settles the modal profile of the normative in our toy state spaces. Against this background, one can evaluate examples that fall on opposing extremes of the irrelevance spectrum. Let me consider the following propositions and their truthmakers:

- (**D**) Sarah promises to donate one percent of her income. = $\{d\}$
- (**N**) Sarah ought to donate one percent of her income. = $\{n\}$
- (**D** \vee **N**) Sarah promises to donate one percent of her income or Sarah ought to donate one percent of her income. = $\{d, n, d \sqcup n\}$

The aim is to make precise the intuition that a transition from D to N is significant to the modal autonomy question, while a transition from D to $D \vee N$ is not significant to the modal autonomy question. So, consider our four modalized toy state spaces: the transition from D to N requires that all states incompatible with N are incompatible with D . That is true within a modalized state space based on $\mathcal{S}_{sC}^\diamond$ or \mathcal{S}_E^\diamond that possess non-minimal degrees of entanglement. It is not true within a modalized state space based on \mathcal{S}_I^\diamond . The reason is that n' which is incompatible with the conclusion N is compatible with the premise D , if one considers \mathcal{S}_I^\diamond . So, only if one constrains the modal profile of the normative like in $\mathcal{S}_{sC}^\diamond$ or \mathcal{S}_E^\diamond , one can construct an intuitively genuine counterexample. This illustrates that N 's normative content is relevant to the entailment relation because we need to make assumptions about the modal profile of N 's normative truthmakers.

To construct a spurious counterexample like the transition from D to $D \vee N$ requires nothing similar. D entails $D \vee N$ no matter which modalized state space one considers, in particular, the modalized state space based on \mathcal{S}_I^\diamond . Recall that \mathcal{S}_I^\diamond represents a scenario where the basic normative states are independent from the descriptive states. Still, one can construct the counterexample. States incompatible with $D \vee N$ are incompatible with each of the disjuncts. Thus, they are incompatible with D . If one considers \mathcal{S}_I^\diamond , the only consistent state that is incompatible with the conclusion is $d' \sqcup n'$. The fact that this state is incompatible with D is just based on assumptions about descriptive states. Thus, even on an account that assumes the minimal degree of entanglement, the spurious counterexample can be constructed. This illustrates that the normative content of $D \vee N$ is irrelevant to the entailment relation because no assumptions about the modal profile of $D \vee N$'s normative truthmakers are required. These considerations show in which sense autonomists need not be worried about spurious examples that fall on one extreme of the irrelevance spectrum. One can construct those examples even though the modal profile of normativity can be reasonably described as independent.

In order to illustrate this result, I have used modalized *toy* state spaces. Note that the toy state spaces I used reduce complexity among two dimensions. On

the one hand, they reduce what one may think of as the range of possibilities, for there is a fairly limited number of normative states in which the world can be. Likewise, there are just two descriptive states in which the world can be. On the other hand, the toy state spaces reduce the complexity of basic normative states, for they do not contain non-trivial descriptive parts.

Increasing complexity on the first dimension, extends the set of normative and descriptive states. Independently of increasing complexity among this dimension, one can construct the discussed kind of spurious counterexamples that do not make any demands on the modal profile of the normative domain. If one would only allow for complexity among this dimension, it would be relatively easy to tell which counterexamples are spurious. As long as one can sharply separate descriptive and basic normative states, one can easily assess whether modal relations are based on assumptions about the modal profile of normativity. As soon as one increases complexity on the second dimension, one faces normative states with non-trivial integral descriptive parts. This drastically increases the complexity of assessing counterexamples. The reason is that normative states with integral descriptive parts are incompatible with all descriptive states incompatible with the integral parts. In order to assess whether the conclusion's normative content is irrelevant, one has to assess whether incompatibilities are just due to descriptive parts or whether they are additionally due to the normative state as a whole. These cases fall in-between the extremes of the irrelevance spectrum. To address them, a fine-grained account of irrelevance is needed. Before turning to an existing account of irrelevance, I shall make a brief aside: it has been argued that the need for an account of irrelevant conclusions is not limited to the autonomy debate. The existing account that I am going to discuss afterwards is developed against the background of issues in other debates that bear some resemblance to disjunctive cases in the autonomy debate. It will be illuminating to be aware of those cases to overcome the worry that modifications of the modal autonomy thesis are *ad hoc*.

Aside: Irrelevant Parts Abound

In the next section, I shall discuss a syntactic account of irrelevant conclusions in Schurz (1991, 1997). Schurz takes his account to solve issues in several debates. One of them is a formal account of theory confirmation. Before presenting issues that arise in that debate, let me explain why being aware of those other debates is helpful. In the autonomy debate, one objection against modified explications of a logical autonomy thesis is that they are *ad hoc* attempts to immunize an autonomy thesis against counterexamples. In other words, suppose there is an extremely complicated sense in which normativity is autonomous. A legitimate question would

be: why should we care? My overall aim in this book is to provide an autonomy thesis that specifies a sense of autonomy that is worth caring about. I take issues in other debates to show that there is an intuitive notion of *consequence* employed in those debates that requires sensitivity to *relevance*. Accordingly, the need for an account of irrelevant conclusions is not *ad hoc*.

Schurz (1991) points out that *prima facie* plausible assumptions about *deductive theory confirmation* trivialize the concept.³ He argues that the trivialization is due to *irrelevant parts* of conclusions. Let me first turn to deductive theory confirmation. Scientific theories can be used to make predictions. If a theory's predictions turn out to be true, this counts in favor of the theory. For example, Relativity theory predicted the existence of black holes. When black holes were discovered, their discovery confirmed Relativity theory. The following definition aims to provide a formal explanation of this practice:

Deductive Theory Confirmation: A sentence S confirms a theory T if (i) S is not necessary, (ii) T is consistent, (iii) S is true, and (iv) $T \vdash S$.⁴

Moreover, the following strengthening principle is taken to be a reasonable requirement on deductive theory confirmation. Suppose that S confirms T . If true S' is consistent with T and entails S , then S' confirms T . Note that the strengthening principle means that with respect to true sentences that are consistent with T the category *T-confirmer* is closed under converse-entailment. The definition of deductive theory confirmation together with the strengthening principle largely trivializes the notion of theory confirmation. To see why, consider a true contingent sentence S that is consistent with a consistent theory T , while $\neg S$ is consistent with $\neg T$. I show that S confirms the theory T . Consider the sentence $T \vee S$. To see that $T \vee S$ confirms the theory T , we go through the four conditions specified by the definition. First, $T \vee S$ is not necessary because we assumed that $\neg S$ is consistent with $\neg T$. Second, T is consistent by assumption. Third, $T \vee S$ is true because S is true by assumption. Fourth, $T \vdash T \vee S$. Thus, $T \vee S$ confirms theory T . By the strengthening principle, S confirms T because it is true and consistent with T and entails $T \vee S$, which I have just shown to be a *T-confirmer*. But this result trivializes any notion of theory confirmation. Every arbitrary true sentence S that is consistent with an arbitrary consistent theory T , while $\neg S$ is consistent with $\neg T$ confirms the theory T . The following example illustrates this counterintuitive result: according to this notion of deductive theory confirmation, that tapirs live in the rain forest confirms Relativity theory.

One might think that in response one should reject the strengthening principle. Yet, even without the strengthening principle, we obtain the

result that $T \vee S$ confirms the theory T . The reason is that $T \vee S$ is a true consequence of T . But $T \vee S$ can be true in virtue of S . Since S might be a disjunctive part of the conclusion that is irrelevant to the fact that T entails $T \vee S$, it seems implausible to assume that $T \vee S$ confirms T . A plausible diagnosis is that only those consequences of a theory that do not contain *irrelevant parts* count as proper consequences of the theory and thereby as putative theory confirmers.

As set out in the beginning, the example illustrates that an account of irrelevant conclusions is not only interesting from the perspective of an autonomy debate. An account that allows us to identify those parts of a conclusion in virtue of which an entailment relation holds may put us in a better position to address issues that result from partially irrelevant conclusions. This concludes my digression on the topic of theory confirmation. Against this background, I shall now consider Schurz's syntactic account of irrelevant conclusions.

Irrelevance and Syntactic Difference-Makers

In this section, I discuss an existing syntactic account of partially irrelevant conclusions developed in Schurz (1991, 1997).⁵ Schurz's account is based on a syntactic criterion of difference-making. I argue that his syntactic focus fails to provide us with a reliable guide to relevance. Against the background of Schurz's syntactic account, I develop an alternative semantic account in the next section. Let me flag that I confine myself to discussing Schurz's account of irrelevant conclusions, while largely setting aside his explication of the logical autonomy thesis. I take his explication of the logical autonomy thesis to be subject to the same objections raised against syntactic taxonomies and conservativeness explications of logical autonomy theses in failing to reveal anything distinctive about normativity.

Schurz develops a general account of partially and completely (ir-)relevant conclusions in terms of uniform predicate substitution.⁶ The account is based on substitution *salva validitate* and thereby makes use of a non-difference-making approach. Through the method of substitution, it is tested whether changes to subformulas or predicates involved in the conclusion make a difference to the inference's validity. If substitution does not change the inference's validity, the substituted parts are irrelevant to the inference. On Schurz' account of uniform substitution of predicates, propositional variables are treated as zero-ary predicates. A complete presentation of Schurz' account would require a number of technical details concerning substitution functions. For my current purposes, the technical implementation is not required and thus might be set aside.⁷ *Partially* and *completely irrelevant conclusions* are defined

in terms of substitutions that do not make a difference to the inference's validity:

Let $\Gamma \vdash C$ be a given deductive inference. The conclusion C is

- *completely irrelevant* iff *all* occurrences of *all* predicates of C
- *partially irrelevant* iff *some* occurrences of *some* predicates of C

are replaceable by any other predicate *salva validitate* of $\Gamma \vdash C$.⁸

Before considering examples, let us turn to the application to normative cases. First, note that the account is based on a syntactic taxonomy. Schurz is concerned with inferences from premises that do not involve deontic operators to conclusions that involve deontic operators. The application of the account to the normative case requires the notion of an O -restricted substitution, where 'O' is the ought-operator. For any substitution function σ and formula A , $\sigma^O A$ denotes the result of performing the O -restricted σ -substitution in A which replaces every predicate F in A by (σF) on exactly those occurrences that lie in the scope of the O -operator (O -occurrences).⁹ To see the difference between an unrestricted and an O -restricted substitution, consider the following example: let $A := P \rightarrow OP$ and $(\sigma P) = Q$. The result of the unrestricted substitution is $\sigma A = Q \rightarrow OQ$, while the result of the O -restricted substitution is $\sigma^O A = P \rightarrow OQ$. Based on the notion of O -restricted substitutions Schurz defines partially and completely O -irrelevant conclusions:

Let $\Gamma \vdash C$ be a given deductive inference. The conclusion C is

- *completely* O -irrelevant iff *all* O -occurrences of *all* predicates of C
- *partially* O -irrelevant iff *some* O -occurrences of *some* predicates of C

are replaceable by any other predicate *salva validitate* of $\Gamma \vdash C$.⁸

It will be instructive to go through some examples. The bold occurrences of predicates are replaceable *salva validitate*. We start with two examples of completely irrelevant conclusions:

- i $P \wedge \neg P \vdash \mathbf{R} \rightarrow \mathbf{OS}$, and
- ii $Q \vdash \mathbf{OP} \vee \neg \mathbf{OP}$.

The inferences are still valid even if *all* predicates of the conclusion are replaced. Two examples that are completely O -irrelevant but not completely irrelevant are:

- i $P \vdash \neg P \rightarrow (\mathbf{OQ} \wedge \mathbf{OR})$, and
- ii $P \rightarrow Q \vdash (Q \rightarrow \mathbf{OR}) \rightarrow (P \rightarrow \mathbf{OR})$.

All predicates within the scope of the ought-operators are replaceable *salva validitate*. A partially but not completely O-irrelevant conclusion is $OP \vdash O(P \vee Q)$. Finally, a partially irrelevant conclusion that is not partially O-irrelevant is $OP \vdash \neg OP \rightarrow Q$. Note that this also illustrates that the list of spurious counterexamples to logical autonomy theses is longer than the list of cases that are commonly discussed.

Schurz provides a detailed and rich discussion of Is-Ought inferences in different alethic-deontic logics without bridge principles.¹⁰ He proves the following generalizable explication of the normative autonomy thesis: premises that do not involve any O-occurrences do not entail a conclusion that involves O-occurrences unless the conclusion is completely O-irrelevant.¹¹ Note that this result holds for *all* propositional operators in a logic without bridge principles for the operator in question.

I shall now explain why a syntactic difference-maker account does not provide a reliable guide toward relevant conclusions. The account tells whether a predicate is necessary for the validity of a particular inference; it does not tell whether the predicate is *relevant* in the context of the inference. The substitution-test identifies a *minimal* truth-value assignment that is required for truth-preservation starting from a set of premises. This suggests that (partial) irrelevance emerges from *redundancy*. By contrast, I argue that, particularly in the context of autonomy, (partial) irrelevance results from *unconnectedness*.

The inference $P \wedge Q \vdash P \vee Q$ illustrates the difference. According to Schurz's account, the conclusion is partially irrelevant (for every instantiation). One can replace Q (P) in the conclusion by any other formula and the inference is still valid. Note that the inference is partially irrelevant with respect to P and partially irrelevant with respect to Q , but it is not partially irrelevant with respect to P and Q together. For concreteness, consider the following instance: 'There are elephants in the zoo and there are turtles in the zoo' entails 'There are elephants in the zoo or there are turtles in the zoo'. If one disjunct is held fixed, the other disjunct is *redundant* with respect to the entailment relation. However, there is a symmetry between both disjuncts because the entailment relation holds in virtue of each of them individually. This is a case of overdetermination. The entailment relation holds in virtue of P , but also in virtue of Q . Both disjuncts are *connected* to the premise. Contrast this with the conclusion 'There are elephants in the zoo or Macron is the French president'. In contrast to the previous, this conclusion is plausibly partially irrelevant, for the second disjunct is *unconnected* to the premise. Intuitively, the second but not the first entailment can be seen to hold for *reasons having nothing to do with* the content of the second disjunct. In the next section, I develop an alternative account of irrelevant conclusions that does justice to a conclusion's redundant, but *connected* content.

Irrelevance and Exact Exclusion

Recall that the conclusion's normative content in a spurious counterexample is *irrelevant* to the entailment relation. The difficulty for any approach to deal with spurious counterexamples is to draw an appropriate line between examples that fall in-between the extremes of the irrelevance spectrum. That is, examples where the normative conclusion is entailed by the descriptive premises in virtue of the conclusion's descriptive content, but where it is hard to tell whether in addition to that the normative content is also relevant. Let me give some potential examples: One such example is the normative conclusion [Mary ought to help or she is permitted not to help] that is entailed by the descriptive premise [Mary exists]. It is difficult to assess whether the normative content is relevant to the entailment relation. Another example concerns disjunctive conclusions with a descriptive and a normative disjunct each of which is entailed by the descriptive premise. For example, suppose that the descriptive premise [God exists] entails the normative conclusion [We ought to praise God or an almighty being exists] where each of the disjuncts is entailed by the premise. Again, in the example the normative content seems to be redundant, but still relevant. These difficulties result from forms of *overdetermination*. A syntactic account of irrelevant conclusions turned out to be inappropriate to account for that difficulty.

The state-based account of irrelevant conclusions that I am going to propose in the following two sections is based on a fine-grained notion of relevance that helps to systematically assess these difficult cases. The starting point for my account is the following observation: according to the definition of entailment within the truthmaker framework, P entails Q iff all states *incompatible* with Q are *incompatible* with P .¹² The definition makes use of the notion of incompatibility on two occasions. My proposal in a nutshell is this: we need a more fine-grained, and in particular, non-monotonic notion of incompatibility to guarantee that the conclusion is not partially irrelevant to the entailment relation. This is why, in this section, I shall develop a notion of *exact incompatibility*, or more precisely, a notion of *exact exclusion*.¹³ Based on this notion, I develop a semantic account of irrelevant conclusions relative to a particular subject matter in the next section. It will take me some time to work out the details of a notion of exact exclusion that can handle difficult cases of overdetermination. Yet, the notion of exact exclusion provides the ground for a fine-grained account of irrelevant conclusions that in turn is essential for a satisfying distinction between spurious and genuine counterexamples. As will become clear, it also provides the basis for an elegant formulation of the modal autonomy thesis at the level of states. But this has to wait until the notion of exact exclusion is at our disposal.

Classical entailment does not require that the relata's contents are connected. Thus, there is room for irrelevant parts. This well-known trait is most vividly illustrated by examples involving necessarily true conclusions or necessarily false premises. To incorporate relevance constraints, a more fine-grained notion of entailment is required. Recall that a *modal* notion of *truthmaking*, according to which P 's truthmakers are the possible worlds that are *compatible* with P , has been criticized for its coarse-grainedness in the context of a descriptive-normative divide. In response, I focused on a notion of *exact truthmaking*, according to which states must be wholly relevant to P 's truth to count as exact truthmakers. A similar strategy provides the ground to distinguish between *inexact* and *exact* exclusion among states. Roughly, for a state s to *exactly exclude* another state t , s as a whole must help to exclude t . Exact exclusion is a relation that is relevance-sensitive, but still modal, in contrast to relations that additionally require forms of, for example, explanatory priority.¹⁴

To begin with, note that two incompatible states exclude each other. However, this notion of exclusion is not exact. To illustrate the inexactness, consider a world-state w and suppose that in w Macron is the French president. Let p be a state to the effect that Macron is the French president. Thus, we have $p \sqsubseteq w$. Now, consider a state p' to the effect that France has no president. w is incompatible with p' and thus excludes p' . However, many of w 's parts do not matter to the fact that w excludes p' . Consider, for example, the state that it is sunny and assume that it is part of w . That it is sunny is completely unrelated to the question of whether France has a president. w and p' are incompatible because w contains a proper part p that is incompatible with p' , while we assume the following principle concerning the set of consistent states S^\diamond :

Downwards- S^\diamond -Closure: If $s \in S^\diamond$ and $t \sqsubseteq s$, then $t \in S^\diamond$.

This principle entails that two states are incompatible with each other as soon as some parts of them are incompatible with each other. Note that this is a plausible assumption: once we obtained an inconsistency, we cannot fix it by merely adding *further* facts. The above example, however, illustrates that inconsistent states such as $p' \sqcup w$ can contain parts that can be seen as non-contributing parts, e.g., that it is sunny.¹⁵ w excludes p' but it contains parts that do not help to exclude p' . In those cases, we may say that w *inexactly excludes* p' . By contrast, the state that Macron is the French president excludes the state that France has no president but it does not involve parts that do not help to exclude p' . In this case, we may say that the state *exactly excludes* the other state. In what follows, I develop and discuss this notion of *exact exclusion*.

Let me again emphasize the similarity with the notion of exact truthmaking. For s to be an exact truthmaker of P , s must be relevant as a whole to P 's truth. Similarly, for s to exactly exclude t , s must be relevant as a whole to the exclusion of t . One might ask whether we can make that notion of exact exclusion formally precise. In light of the world-president example, it seems natural to define exact exclusion in terms of incompatibility and parthood. Consider the following definition: s exactly excludes t iff s and t are incompatible and no proper part of s is incompatible with t . The proposal yields the desired result for the world-president example: w and p' are incompatible, but w contains a proper part p that is incompatible with p' . Thus, w does not exactly exclude p' .

Note that this first attempt to define exact exclusion requires that s is *minimal* with respect to its incompatibility with t . However, we already encountered a reason not to adopt a minimality requirement. Consider a state to the effect that Macron is the French president and it is sunny ($p \sqcup s$). This state is incompatible with the state that France has no president and it is rainy ($p' \sqcup s'$). Since $p \sqcup s$ contains p which is incompatible with $p' \sqcup s'$, $p \sqcup s$ is not minimal with respect to its incompatibility with $p' \sqcup s'$. Accordingly, the state that Macron is the French president and it is sunny would not count as being relevant as a whole to the exclusion of the state that France has no president and it is rainy. Yet, all parts of the former state intuitively help to exclude the latter state.¹⁶

One might wonder whether the definition can be fixed by adding a further clause. A state s that contains a proper part v that is incompatible with t still exactly excludes t iff s without v is also incompatible with t . Roughly, it is required that we can separate s into parts each of which is incompatible with t . However, cases that are based on infinite sets arguably generate some trouble for this proposal.¹⁷ Fortunately, in the end, I can stay neutral with respect to these cases, but let me briefly sketch the issue. Suppose there is an infinite set of states $\Sigma \subseteq S$. Let us stipulate that in the state space there is a state f such that f is incompatible with every infinite fusion of states that belong to Σ , but compatible with every finite fusion of states in Σ . It might be difficult to give an example of the form *the state that ...* Note that what might seem to be natural candidates like the state that there are finitely many natural numbers have the form of existential quantifications which makes them probably too indeterminate to count as state in the state space. Yet, f 's structural features suffice to illustrate a putative issue of the definition under consideration. As defined, f is not exactly excluded by any infinite fusion of states in Σ because infinite fusions contain finite fusions as proper parts that are compatible with f . Intuitively, however, infinite states should at least in principle be able to exactly exclude f iff all the parts of the infinite state help to exclude f .

Fortunately, I can sidestep this issue. The outlined conditions in the definition might not be necessary, but they are plausibly sufficient for s to exactly exclude t . For the current purposes, I shall assume that exact exclusion is a primitive, but constrained notion in the framework. Let me restate the constraints: States that are minimally incompatible with a state t exactly exclude t . I call these states t 's *basic excluders*. Additionally, I assume that any fusion $s = s_1 \sqcup s_2 \sqcup \dots$ where each s_i exactly excludes t exactly excludes t .

There is another worry that I would like to address before finally turning to some special cases of exact exclusion. One might be worried that the notion of exact exclusion relies on a notion of relevance. According to the informal gloss, states that exactly exclude t are wholly relevant to the exclusion of t . Since our aim is to define a notion of irrelevant conclusions, one might object that any such definition would be circular. However, as will become clear in due course, many counterexamples to autonomy theses can be analyzed in terms of basic excluders and fusions of basic excluders. In those cases, we do not presuppose any unanalyzed relevance considerations. With respect to the other cases, it is worth pointing out that I do not aim to give a reductive analysis of relevant conclusions. Even if we have to rely on some primitive notion of relevance, a state-based analysis will help us to clarify relevant features of complex cases.

Let me now turn to some special cases of exact exclusion concerning *inconsistent* states. Some counterexamples to a simple logical autonomy thesis involve necessarily true conclusions or necessarily false premises. Accordingly, I briefly discuss two cases: (i) consistent s exactly excludes *inconsistent* t and (ii) *inconsistent* s exactly excludes consistent t .

First, suppose that t is inconsistent. Hence, t has only a single basic excluder, namely the null-state. The minimal state that is incompatible with an inconsistent state is the null-state. If s is not the null-state, then s is neither a basic excluder nor the fusion of basic excluders of t . Consider the following plausible principle: if a state u exactly excludes v and $v \sqsubseteq w$, then u exactly excludes w . Accordingly, basic excluders of consistent parts of t exactly exclude t . So, s exactly excludes inconsistent t iff it is the null-state or a basic excluder of a consistent part of t . Second, suppose s is inconsistent. Suppose further that no proper part of s is incompatible with t . Consequently, s is minimally incompatible with t and thereby a basic excluder of t . This yields somewhat unintuitive results. For example, the inconsistent state to the effect that Macron is the French president and France has no president exactly excludes that it is sunny. In response, we shall restrict exact exclusion so that only consistent states exactly exclude other states.

To sum up, exact exclusion is a non-monotonic, non-symmetric exclusion relation. That is, if we add an arbitrary state to an exact excluder of

a state s the result is not guaranteed to be an exact excluder; moreover, an exact excluder of a state s is not guaranteed to be exactly excluded by s . I shall now turn to an account of irrelevant conclusions in terms of exact exclusion that in turn provides the ground for my modal autonomy thesis.

A State-Based Account of Irrelevant Conclusions

It is important to bear in mind that I eventually aim to give an account of spurious counterexamples. On the account that I am going to propose, the conclusion of spurious counterexamples is such that the conclusion's normative content is irrelevant to the entailment relation. However, it is sometimes difficult to tell whether it is just the conclusion's descriptive content that is relevant or whether its normative content is also relevant. My account of irrelevant conclusions is based on the idea that states incompatible with a partially irrelevant conclusion bear a merely inexact relation to the conclusion or the premises. To make this idea precise, I employ the outlined notion of exact exclusion. The variety of putative counterexamples to a modal autonomy thesis results from logical complexity at the propositional level. Descriptive and normative content can be entangled in complex configurations. My preferred formulation of the modal autonomy thesis is straightforward: no descriptive premises entail a normative conclusion unless the conclusion is entailed in a normatively irrelevant way. However, the formal account of normative, and in general, subject-matter-based irrelevance will require some complex conditions in order to account for different ways in which irrelevance may arise. Once I have worked out these complex conditions, it will become clear that the autonomy thesis is equivalent to a simple and elegant formulation at the level of states. The state-based account thus reduces complexity at the propositional level, which greatly facilitates the assessment of complex cases that fall in-between the relevance spectrum.

Varieties of Irrelevant Conclusions

I now turn to partially irrelevant conclusions. Suppose a set of premises Γ entails a conclusion C . Within the state-based framework, this means that every state s incompatible with C is incompatible with Γ . Irrelevance can emerge due to each of the two incompatibility requirements because incompatibility is understood in terms of *inexact* exclusion. To guarantee that the conclusion is wholly relevant to the entailment relation I need a refined account of entailment in terms of *exact* exclusion.¹⁸ Let me illustrate how irrelevance can emerge due to each of the two incompatibility requirements.

First, consider an example where states incompatible with the conclusion bear a merely inexact relation to the *premises*. Consider $P \models P \vee Q$ where P and Q are unrelated. States incompatible with $P \vee Q$ must be incompatible with both disjuncts. Let p' be incompatible with P and q' be incompatible with Q . It is easily verified that $p' \sqcup q'$ fails to exactly exclude any $p \Vdash P$ (truthmakers of the premise) because q' does not help to exclude any p . This motivates an additional notion of exclusion: I shall say that a *state* s exactly excludes a *proposition* P iff s is the fusion of exact excluders for every $p \Vdash P$.¹⁹ Accordingly, $p' \sqcup q'$ fails to exactly exclude P (the premise).

Second, consider an example where the *conclusion* bears a merely inexact relation to states that are incompatible with it. Consider $P \models P \vee (P \wedge Q)$ where P and Q are unrelated. To be incompatible with P suffices to be incompatible with the conclusion. These states bear an exact relation to the premise. Yet, not all truthmakers of the *conclusion* bear an exact relation to those states. Suppose p' exactly excludes P . It is thus incompatible with the conclusion. However, some states that make the conclusion true such as $p \sqcup q$ do not exactly exclude p' . Parts of the conclusion are irrelevant because states that make these parts true do not help to exclude states incompatible with the conclusion. Accordingly, I define partially irrelevant conclusions as follows:

Partial Irrelevance: Let Γ entail C . C is a partially irrelevant conclusion iff all consistent t incompatible with C either i) fail to exactly exclude Γ or ii) some $c \Vdash C$ does not exactly exclude t .

To get a better grip on the definition, one may think of a partially irrelevant conclusion C as follows: there is a set of states incompatible with C . These states split up into two camps (one of which might be empty). On the one hand, there are states that are *oversized* which means that they do not exactly exclude the premises because they contain non-contributing parts (see Figure 5.1, left part). On the other hand, we have states that are *undersized* which means that truthmakers of the conclusion do not exactly exclude them because the truthmakers of the conclusion contain non-contributing parts (see Figure 5.1, right part). Figure 5.1 illustrates both sources of partial irrelevance with respect to the example $A \models A \vee (A \wedge B)$ where the arrows represent exact exclusion.

The outlined notion of partially irrelevant conclusions yields the desired results with respect to cases of overdetermination like in the zoo example. Recall that in cases where the entailment relation holds in virtue of both disjuncts of the conclusion individually, the conclusion is intuitively wholly relevant. In cases where it holds merely in virtue of one disjunct, the conclusion is partially irrelevant. On the present account, states that

OVERSIZED STATE

UNDERSIZED STATE




Figure 5.1 Partial irrelevance and non-contributing parts.

are incompatible with $P \vee Q$ such as $p' \sqcup q'$ and states that are incompatible with $P \vee R$ such as $p' \sqcup r'$ both *inexactly* exclude $P \wedge Q$. However, only $p' \sqcup q'$ *exactly* excludes $P \wedge Q$ which means that $P \vee R$ but not $P \vee Q$ is partially irrelevant to the entailment relation.

Note that for a conclusion to be partially irrelevant it suffices that some arbitrary part of the conclusion is irrelevant to the entailment relation. For the purposes of *autonomy theses*, we need to make sure that not *all* parts of a conclusion that are about a particular *subject matter* are irrelevant. In the normative case, we need to make sure that not all normative parts are irrelevant.²⁰ If we are interested in the hedgehog domain, we need to make sure that not all hedgehog-parts are irrelevant. In order to cover the cases where all of the conclusion's subject-matter-related parts are irrelevant, we need to focus on particular states. The following state-based analysis of normatively irrelevant conclusions makes sure that the conclusion's normative content is irrelevant to the entailment. Let Γ entail normative C :

Normative Irrelevance: C is a normatively irrelevant conclusion iff for all states t that exactly exclude C either i) no normative part²¹ $n \sqsubseteq t$ exactly excludes Γ or ii) no normative $c_N \sqsubseteq t$ exactly excludes $n \sqsubseteq t$ where n exactly excludes Γ .

It is worth mentioning that the label *normatively irrelevant conclusion* might not be the best way to conceive of those conclusions. What is wrong with spurious counterexamples is that the conclusion's normative content is irrelevant to the *entailment relation*. This does not mean that the conclusion is in any general sense normatively irrelevant. However, for ease of presentation, I shall continue to use that label.

Let me consider examples. Consider $D \vDash D \vee N$ where D and N are unrelated. States that exactly exclude the conclusion are fusions of states

that exactly exclude the disjuncts. Let $d' \sqcup n'$ be such a state. Neither $d' \sqcup n'$ nor any normative part of $d' \sqcup n'$ exactly excludes D . To see why the second clause is required, suppose some normative state m' exactly excludes D . In that case, $m' \sqcup n'$ exactly excludes the conclusion and some of its normative parts, namely m' , exactly excludes the premise. However, m' is unrelated to the conclusions normative content. The second condition accommodates for this case because the normative truthmakers of the conclusion n and $d \sqcup n$ do not exactly exclude m' . Thus, the conclusion is normatively irrelevant. Note that if D entails N , then the account yields the desired result that the conclusion is normatively relevant. Second, consider $D \models D \vee (D \wedge N)$ where D and N are unrelated. States that exactly exclude the conclusion always contain states that exactly exclude D . Normative states that exactly exclude the conclusion such as $d' \sqcup n'$ do not exactly exclude D ; neither do its normative parts. Descriptive states that exactly exclude the conclusion such as d' vacuously satisfy the first condition. Again, if there is a normative m' that exactly excludes D , then it is not exactly excluded by $D \vee (D \wedge N)$'s normative truthmaker $d \sqcup n$. The outlined definition can be generalized to an arbitrary subject matter SM .

SM-Irrelevance: C is a SM -irrelevant conclusion iff for all states t that exactly exclude C either i) no SM -part $v \sqsubseteq t$ exactly excludes Γ or ii) no SM -state $c_{SM} \Vdash C$ exactly excludes $v \sqsubseteq t$ where v exactly excludes Γ .

On the syntactic account, arguments with necessarily false premises such as $P \wedge \neg P \models Q$ and arguments with necessarily true conclusions such as $P \models Q \vee \neg Q$ have completely irrelevant conclusions. These examples are already covered by the subject-matter-based irrelevance account. For concreteness, I consider the normative case, but the reasoning applies in general. First, consider $D \wedge \neg D \models N$. For N to be normatively relevant, a state that exactly excludes N must contain normative parts that exactly exclude $D \wedge \neg D$. This is only the case if that normative state exactly excludes D or $\neg D$. Normative parts of states that exactly exclude N do not exactly exclude D or $\neg D$ unless N 's normative content is related to one of them. If D entails normatively relevant N , then normative states that exactly exclude N exactly exclude D . Hence, the conclusion of $D \wedge \neg D \models N$ is normatively relevant in those cases. If N is entailed by neither conjunct, the conclusion is normatively irrelevant. Note that on the syntactic account the conclusion of $N \wedge \neg N \models N$ is normatively irrelevant, while according to the present account it is normatively relevant.

Finally, consider necessarily true conclusions such as in $D \models N \vee \neg N$. Note that states incompatible with the conclusion are inconsistent. Thus,

no state exactly excludes the conclusion and the definition of a normatively irrelevant conclusion is vacuously satisfied. It might be seen as an implausible consequence of the definition that $N \models N \vee \neg N$ has a normatively irrelevant conclusion. More importantly, I have argued that if D entails N , the conclusion of $D \models N \vee \neg N$ should count as normatively relevant. To account for this assessment, I adopt the following clause for cases that vacuously satisfy the definition: no normative state that exactly excludes one of C 's exact truthmakers exactly excludes Γ . It is easily verified that according to this clause, the inference $D \models N \vee \neg N$ has a normatively irrelevant conclusion iff D entails neither normatively relevant N nor $\neg N$.

This concludes my account of irrelevant conclusions. Before turning to my preferred explication of the modal autonomy thesis, I shall address a question that might arise. The syntactic account identifies predicates that are irrelevant to the inference; the Finean Gap Principle tells us that only $(\sim Q)_D$ is relevant to the entailment relation. So, one might ask whether the outlined account tells us which *propositional parts* are irrelevant.

Irrelevant Propositional Parts

To address that question, I must specify the relevant notion of propositional part. We already encountered two notions of propositional part, namely *disjunctive* and *conjunctive* parts. My first aim is to show that these notions of propositional part do not cover all examples of irrelevant conclusions. Instead, I shall consider another related notion of propositional part discussed in Krämer (2017). I argue that it is a special trait of the normative domain that under plausible assumptions about normativity, some normative conclusions are normatively irrelevant, yet not in virtue of particular propositional parts.

First, consider conjunctive and disjunctive parts. At first glance, disjunctive parts seem to be promising candidates because in cases like $D \models D \vee N$, N is a disjunctive normative part of the conclusion that is irrelevant to the entailment relation. Yet, not all examples have irrelevant normative disjunctive parts. Reconsider $D \models D \vee (D \wedge N)$. Intuitively, as well as according to the formal account, N is no disjunctive part of $D \vee (D \wedge N)$ because N is no subset of $D \vee (D \wedge N)$ on the truthmaker-based account. N is no conjunctive part neither because descriptive states that make the conclusion true do not contain N 's normative truthmakers thereby violating one of the conditions for a conjunctive part.

Accounting for examples such as $D \models D \vee (D \wedge N)$ requires a notion of a conjunctive-disjunct-part (*cd-part* for short) as it is employed in Krämer (2017). A *cd-part* is a conjunctive part of at least one disjunctive part. In the example, N is a conjunctive part of the disjunctive part $(D \wedge N)$ of the conclusion. Thus, we have an irrelevant normative *cd-part*. Since every

proposition is a conjunctive part of itself, every disjunctive part is also a *cd*-part. Accordingly, one might suggest that a conclusion is *SM*-irrelevant iff no *SM-cd*-part makes a difference to the entailment relation.²²

In its application to the normative domain this proposal comes to its limits. Recall that normative states can contain *integral* descriptive parts. If a normative state *n* contains a descriptive integral part *d*, then *d* cannot be supplemented by a proper part of *n* to obtain *n*. This is, the remainder cannot be represented by a state in the state space. Consider, for example, a normative state to the effect that Sarah ought to donate. This state contains the descriptive state that Sarah exists. The descriptive state is an integral part of Sarah's obligation to donate because there does not seem to be any normative state other than Sarah's obligation that could be added to Sarah's existence to obtain Sarah's obligation. We considered descriptive integral parts to avoid any commitment to purely normative states. I have argued that this is an important strength of the account.

Against this background, reconsider the following putative counterexample to an autonomy thesis: [Mary exists] entails [Mary ought to attend the meeting or Mary is permitted not to attend the meeting].²³ For the current purposes, let us preliminarily accept that the conclusion is normatively irrelevant. I will more carefully discuss and defend this claim in the next section. A first thing to note is that the disjunctive parts of the conclusion are not irrelevant. If we replace one of the disjuncts, the entailment relation is not guaranteed to hold. But what is about other normative *cd*-parts? Any *cd*-part must be a conjunctive part of one of the two disjuncts. But even if there are such parts (e.g., because obligations conceptually contain permissions) they are not irrelevant for the entailment relation. To see why, note that Mary's existence is an integral part of both disjuncts. Each normative state that is part of a truthmaker of the conclusion contains Mary's existence. This means that one cannot separate the conclusion into a *relevant descriptive* and an *irrelevant normative* propositional part. Suppose one would have allowed for purely normative states such that one could separate, for example, Mary's obligation into her existence and a purely normative state that does not contain her existence. In that case, we would be able to identify irrelevant normative *cd*-parts that correspond to the relevant purely normative states. In general, if we are concerned with domains that do not require integral parts we have irrelevant propositional *cd*-parts. The assumption of purely normative states, however, is problematic. Thus, we obtain normatively irrelevant conclusions without irrelevant normative propositional parts.

This concludes my discussion of irrelevant conclusions.²⁴ Some putative counterexamples to autonomy theses emerge because the entailment relation holds in virtue of descriptive content. The challenge for an autonomy debate is to assess whether these cases always generate spurious

counterexamples. I have argued that this is only the case if the conclusion is normatively irrelevant. Since a formally precise account of the relevant notion of irrelevance is now at our disposal, I can return to the assessment of cases falling in-between the extremes along the irrelevance spectrum.

A Relevance-Based Modal Gap

In this final section, I lay down my preferred explication of the logical autonomy thesis: the *Relevant Gap Principle* (*Re-Gap* for short). I show that it can be generalized to arbitrary subject matters, yet is still a substantial claim about normativity. I show that the Relevant Gap Principle is equivalent to a simple and elegant claim about states within the state space. I demonstrate that the account entails satisfying distinctions between spurious and genuine counterexamples by considering cases that fall in-between the extremes of the irrelevance spectrum. To that end, I compare the Relevant Gap Principle with the Finean Gap Principle.

The Relevant Gap Principle

Spurious counterexamples threaten a simple logical autonomy thesis according to which no descriptive premises entail a normative conclusion. Previous considerations revealed that spurious counterexamples to logical autonomy theses involve normative content that is *irrelevant* to the entailment relation. This suggests the following relevance-based explication of a modal autonomy thesis concerning the normative domain and its generalization to an arbitrary subject matter *SM*:

Relevant Gap Principle:	No descriptive premises entail a normative conclusion unless the conclusion is normatively irrelevant. ²⁵
General Relevant Gap Principle:	No non- <i>SM</i> -premises entail a <i>SM</i> -conclusion unless the conclusion is <i>SM</i> -irrelevant.

In the previous section, I developed an account of the relevant notion of irrelevant conclusions. It turned out that in order to accommodate for all cases the account requires some complexity. According to the state-based analysis, a conclusion is *SM*-irrelevant iff for all states *t* that exactly exclude the conclusion either no *SM*-parts of *t* exactly exclude the premises, or no *SM*-states that make the conclusion true exactly exclude *t*'s *SM*-parts that exactly exclude the premises. Admittedly, the complexity is a somewhat unappealing feature of the account. In what follows, I show that the complex account at the propositional level has a simple and elegant equivalent at the state level. The truthmaker-based account can be seen to *reduce* logical complexity at the propositional level which in turn

provides the ground for an assessment of complex counterexamples that are otherwise difficult to assess.

Before turning to the state level, let me briefly highlight two features of the (General) Re-Gap Principle that strike me as important. First, the explication reveals that autonomy requires *relevance* constraints. Like syntactic accounts that focus on non-difference-making occurrences of predicates, the account makes transparent that irrelevance-based entailment relations generate trouble for *any* autonomy thesis. Accordingly, normative irrelevance is just one instance of a general phenomenon.²⁶ At the same time, and this is the second feature, the truth of the logical autonomy thesis, under this explication, depends on features of normativity. To decide whether a conclusion is irrelevant we have to scrutinize the modal profile of normative states. Thus, the (General) Re-Gap Principle combines a *general* strategy to avoid and explain spurious counterexamples with a substantive claim about the normative domain. Let me now state the condition at the state level that all and only state spaces satisfy in which the Re-Gap Principle holds:

State Gap: Normative states do not exactly exclude descriptive states.

To see the equivalence, suppose first that the Re-Gap Principle is false. Thus, some descriptive D entails a normatively relevant N . For this to be the case, some t that exactly excludes N must be such that some normative $n \sqsubseteq t$ exactly excludes D . This means that n is the fusion of exact excluders for each $d \Vdash D$. Given our assumptions about normative states, n cannot be the fusion of descriptive states. So, some normative state exactly excludes some $d \Vdash D$. Thus, the State Gap is false.

Suppose that the State Gap is false. Thus, some normative state n exactly excludes a descriptive state d . Consider $D = \{d\}$ and $N = \{n\}$. We assume that one of N and $\neg N$ must be true. Since D is incompatible with N , D entails $\neg N$. Since $\neg N$ has a normative falsemaker, $\neg N$ is normative. Thus, descriptive D entails normative $\neg N$. Since n is a falsemaker of $\neg N$, it exactly excludes $\neg N$. By assumption, n exactly excludes $D = \{d\}$. Moreover, $\neg N$'s truthmakers must be wholly relevant to the exclusion of N 's sole truthmaker n . Thus, $\neg N$'s is a normatively relevant conclusion and the Re-Gap Principle is false. Let me finally discuss how the Re-Gap Principle helps to assess examples that fall in-between the extremes of the irrelevance spectrum. I do so by comparing the consequences of the Re-Gap Principle with the consequences of the Finean Gap Principle.

Is Everything Fine?

Let me first remind the reader of the Finean Gap Principle: No descriptive proposition P entails a normative proposition Q unless $(\sim Q)_D$ entails $\sim P$,

i.e., unless the descriptive component of every possible state incompatible with Q is incompatible with P .²⁷ The motivating thought is that cases where $\sim Q$'s descriptive content entails $\sim P$ are spurious counterexamples because 'the entailment from P to Q can be seen to hold for reasons having nothing to do with the normative content of Q '.²⁸ I have objected to the approach that in some cases where $\sim Q$'s descriptive content entails $\sim P$, Q 's normative content is also relevant. In those cases, the entailment from P to Q holds for reasons having to do with Q 's normative content. Assessing these cases requires a fine-grained notion of irrelevance.

Recall one of the examples I have discussed. Suppose descriptive D entails normatively relevant N . Thus, we face a genuine counterexample to the autonomy of the normative domain. I have argued that the status as a genuine counterexample transmits to examples like as $D \models D' \vee N$ where D entails D' . After all, D entails $D' \vee N$ in virtue of normative N . The Re-Gap Principle entails this result. States that exactly exclude the conclusion contain normative states that exactly exclude N which in turn exactly exclude D , while $n \Vdash D' \vee N$ exactly excludes this state. Thus, we obtain the desired result that the example is a genuine counterexample.

Let us turn to other difficult cases on the relevance spectrum. We can describe the difficult cases as follows: the entailment relation holds in virtue of the conclusion's descriptive content, but it is hard to tell whether its normative content is also relevant. Recall that only if the normative content is irrelevant, we face a spurious counterexample. I shall discuss two cases that differ with respect to the outcome. The first example that is difficult to assess is the following one: Mary does not exist. Thus, it is not the case that Mary ought to help. Both accounts, the Re-Gap Principle and the Finean Gap Principle, classify the example as a spurious counterexample. That is, the normative content is irrelevant. On the Finean approach, the incompatibility between a state to the effect that Mary ought to help and a state to the effect that Mary does not exist is attributable to a descriptive component of the normative state, namely to Mary's existence. Thus, the entailment relation holds in virtue of the conclusion's descriptive content. On the Re-Gap approach, the crucial question is whether the normative state that Mary ought to help, o_{help} , exactly excludes the descriptive state that Mary does not exist.

The answer is that o_{help} does not exactly exclude Mary's non-existence. Let me explain why. A first thing to note is that o_{help} is no basic excluder of Mary's non-existence because o_{help} contains a descriptive part, namely Mary's existence, that exactly excludes Mary's non-existence. It is no fusion of basic excluders neither because Mary's existence is an integral and thus unsupplemented part of o_{help} . One might think that o_{help} nevertheless exactly excludes Mary's non-existence because it does not contain any non-contributing states. However, there is a reason to resist that result.

It is uncontroversial that there is a difference between Mary's existence and her obligation to help, though it is assumed that the difference is not representable by another state. Whatever the difference is, it does not contribute to the exclusion of Mary's non-existence. Thus, O_{help} fails to *exactly* exclude Mary's non-existence. To see why, compare the following *opposing* states: Mary's obligation to help, O_{help} , and her permission not to help, p_{-help} . The normative contribution of each of these states is just the opposite of the normative contribution of the other one. Yet, both of them are incompatible with Mary's non-existence and there is a symmetry between both cases. O_{help} and p_{-help} exclude Mary's non-existence for exactly the same reason. This suggests that it is merely their descriptive common part that contributes to the exclusion of Mary's non-existence. Accordingly, O_{help} merely inexactly excludes Mary's non-existence. Thus, the example is classified as a spurious counterexample.

The second example has a similar structure, yet a different outcome. Suppose that [Sarah promises to pay five dollars] entails [It is not the case that Sarah is permitted to keep the money]. Intuitively, this example is a genuine counterexample (if it were valid) and thus should falsify an autonomy thesis. On the Finean approach, the assessment depends on whether the incompatibility between a permission to keep the money (a member of $\sim Q$) and a promise to pay is attributable to a descriptive component. Permissions need not have any substantial *descriptive* parts. In that case, the incompatibility is *not* attributable to descriptive states. Accordingly, the example refutes autonomy as explicated by the Finean Gap Principle. However, let me remind the reader that the assumption that permissions do not contain substantial descriptive parts is *necessary* to obtain that result. I have argued that this is an implausible consequence of the Finean account. The account excludes that we conceive of the example as a genuine counterexample, while taking permissions to have descriptive constituents, e.g., the absence of promises to do otherwise.²⁹ On the Finean approach, the outlined mereological profile of normative states is incompatible with denying the autonomy of normativity based on these kinds of examples.

Let me now explain why the Re-Gap Principle is advantageous in those cases. Instead of being committed to a particular mereological profile of normative states, namely that normative states do not contain substantial descriptive parts, an opponent of autonomy is just committed to the *role* played by the normative content. She must defend that the normative content is *relevant* to the entailment. To see why, note that on the Re-Gap approach, we need to settle whether a permission to keep the money exactly excludes a promise to pay. Suppose we take permissions to contain substantial descriptive parts that are incompatible with promises to pay (e.g., their absence). Still, we are free to argue that permission-states as a whole are relevant to the exclusion of promises. To see why permissions

might be wholly relevant, compare a permission state that contains the absence of promises on the one hand and the absence of promises together with the state that 2 is prime on the other hand. In the latter case, we clearly lose the exactness of the exclusion relation. However, the permission state seems to be different; it does not seem to cancel exactness in the same way. The Re-Gap approach thereby makes room for two intuitively compatible assumptions: the denial of autonomy and the claim that normative states have substantial descriptive parts. The account thereby improves the distinction between genuine and spurious counterexamples. That said, the Re-Gap Principle is closely related to the Finean Gap Principle. Autonomy as explicated by the Re-Gap Principle entails autonomy as explicated by the Finean Gap Principle given plausible assumptions about exact excluders.³⁰

This concludes my discussion of the Re-Gap Principle. On the one hand, the Re-Gap Principle explains spurious counterexamples in terms of a generally applicable account of irrelevant conclusions relative to a subject matter; on the other hand, it is sensitive to putative genuine counterexamples. This shows that the Re-Gap Principle is a substantial thesis about the modal profile of normativity.

Summary and Prospects

In this chapter, I have shown that some spurious counterexamples hold even if basic normative and descriptive states are modally independent. This result substantiates the intuitive assessment that spurious counterexamples do not refute normative autonomy. I have argued that putative counterexamples to a modal autonomy thesis fall along a spectrum: some of them are such that the conclusion's normative content is uncontroversially relevant, while others are such that the conclusion's normative content is uncontroversially irrelevant. The difficult cases are those that fall in-between the extremes. In response, I developed a state-based account of irrelevant conclusions. I have argued that irrelevant conclusions result from merely *inexact* exclusion among states. My analysis of irrelevant conclusions relative to a subject matter is thus based on a notion of *exact* exclusion. My account provides a fine-grained distinction between cases where the entailment holds just in virtue of its subject-matter-unrelated content and those cases where it holds in virtue of its subject-matter-related *and* its subject-matter-unrelated content.

Based on my account of partially irrelevant conclusions, I developed the General Re-Gap Principle. According to its normative instance, the Re-Gap Principle, no descriptive premises entail a normatively relevant conclusion. The Re-Gap is based on a general strategy to avoid counterexamples, but eventually depends on distinctive features of the normative

domain. Autonomy thus explicated can be captured at the state level by a simple and elegant claim: normative states do not exactly exclude descriptive states.

The Re-Gap Principle is compatible with some forms of descriptive-normative entanglements. For example, it is compatible with the claim that the existence of an agent is a necessary condition for obligations, permissions, and other normative assessments concerning the agent. I have argued, however, that a satisfying autonomy thesis is not compatible with all kinds of descriptive-normative entanglements. If, for example, promises entail obligations normativity fails to be autonomous. On my account, this result does not depend on further assumptions about the mereological profile of normative states. In the next chapter, I lay out and discuss the possibility that normative debates assume descriptive-normative entanglements that refute the Re-Gap Principle. I argue that there is still room to defend a modal autonomy thesis concerning a particular kind of normativity, namely authoritative normativity. My main claim is that even if we have to restrict the Re-Gap Principle, there is a way to retain a modal autonomy thesis that qualifies as substantial normative autonomy thesis.

Notes

- 1 For the notion of a basic state, see [Fine \(2021, p. 909\)](#).
- 2 This is a legitimate assumption in the current context. I want to show that spurious counterexamples hold even given strong assumptions *in support* of the modal independence of normativity.
- 3 Another example in [Schurz \(1991\)](#) concerns verisimilitude as it is discussed in [Popper \(1962\)](#).
- 4 For this particular formulation, see [Schurz \(1991, p. 397\)](#). See also [Hesse \(1970\)](#).
- 5 For reasons of space, I shall concentrate on Schurz's account. Another approach to deal with partially irrelevant conclusions is the account of *content containment* in [Gemes \(1994, 1997\)](#). Notice that in the previous chapter, we saw that we can formulate a valid autonomy thesis based on analytic entailment that is defined in terms of content containment. I have discussed the hyperintensional account of content containment in [Fine \(2017a\)](#). Gemes's account of partial content is an intensional account in which intensionally equivalent propositions have the same content parts. I have already argued for the importance of hyperintensional distinctions in the context of an autonomy debate. Moreover, the reasons not to be satisfied with an autonomy thesis in terms of analytic entailment arguably transmit to Gemes's account.
- 6 Note that the notion of vacuous occurrences in [Pigden \(1989\)](#) is also based on substitution. In contrast to Pigden's account, Schurz's account accommodates for deontic operators as logical constants.
- 7 [Schurz \(1997, p. §2.4.3\)](#) provides an elaborated account of substitution that prevents confusion of variables.
- 8 [Schurz \(1997, p. 82\)](#).
- 9 [Schurz \(1997, p. 78\)](#).

- 10 Bridge principles are defined as mixed axiom schemata in which at least one schematic letter occurs within and out of the scope of some ought-operator, e.g., $OA \rightarrow \Diamond A$. See Schurz (1997, p. 91).
- 11 Note that the account must respond to counterexamples like (Truth-Teller) by classifying the premise ‘Alfie said that it is obligatory to help’ as normative, for the conclusion is completely normatively relevant. I have argued that this classification is implausible. This assessment gets additional support from Schurz’s explanation why mixed sentences should be classified as normative: mixed sentences play an important role in ethical theories (Schurz 1997, p. 11). However, ethical theories rarely seem to be concerned with Alfie’s utterances. Note that God might be an exception, which in turn seems to be the rule.
- 12 Recall that a state s is incompatible with a proposition P iff s is incompatible with all of P ’s truthmakers.
- 13 The notion of exclusion was introduced in Fine (2017a). Fine uses this notion for what I will call ‘exact exclusion’. I shall keep the exact–inexact-distinction because it seems fairly natural to say that a world with a French president excludes a state to the effect that France has no president, even though—as will become clear in a moment—it does not *exactly* exclude this state.
- 14 See, for example, Krämer (2021). I shall be concerned with priority relations among states in Chapter 8.
- 15 Notice that I use the notion of (non-)contributing parts to elucidate the relation of exact exclusion. As will become clear soon, I take exact exclusion to be a primitive relation that is not defined in more fundamental terms. Difficulties to *define* the notion are closely related to similar difficulties in the context of the debate about proportionality without minimality constraints, see Yablo (2025). Yablo provides a promising account of proportionality that does not invoke minimality constraints. For my purposes, an intuitive grasp on the notion of contributing parts suffices. However, in what follows, I shall introduce constraints on a notion of exact exclusion.
- 16 For an insightful discussion of cases where something makes a relevant contribution without being part of a *minimally* sufficient base, see Yablo (2025).
- 17 For similar examples, see again Yablo (2025). Indeed, some of Yablo’s examples arguably instantiate the outlined pattern.
- 18 Since we engage in a notion of exact exclusion that is characterized by relevance constraints, a natural question is how a logic that corresponds to that relation would look like and how it relates to familiar relevance-based logics. However, I lack the space to discuss these questions here.
- 19 The definition naturally extends to a set of propositions Γ : s exactly excludes Γ iff s exactly excludes $\bigwedge_{P \in \Gamma} P$.
- 20 Note that on the syntactic account one distinguishes between partially and completely normatively irrelevant conclusions. Partially normatively irrelevant conclusions have normative parts that are irrelevant, but may also have a normatively relevant part, e.g., $N_1 \vDash N_1 \vee N_2$. So, we shall say that C is a partially normatively irrelevant conclusion iff all consistent states t incompatible with C either (i) fail to exactly exclude Γ due to a normative part of t or (ii) at least one normative $c_n \Vdash C$ does not exactly exclude t .
- 21 One might wonder why I did not require that all normative states that exactly exclude C do not exactly exclude the premises. Yet, $N \vDash N \vee D$ turns out to be normatively irrelevant according to this definition, while intuitively it is normatively relevant. The proposed account entails that it is normatively relevant.

- 22 One shortcoming of this proposal is that it cannot account for cases of over-determination. This is why my account is not spelled out in terms of difference-making.
- 23 The example originates from [Fine \(2021, p. 898\)](#). Another example to illustrate the same result is $a = b \models O(Fa) \rightarrow O(Fb)$.
- 24 The notion of normatively (ir-)relevant conclusions is not limited to arguments with descriptive premises. It is generally applicable to all arguments.
- 25 For brevity, I shall often say that no descriptive premises entail a normatively relevant conclusion.
- 26 It has already been shown that accounts such as the one in [Russell and Restall \(2010\)](#) provide a unified treatment of several inference barriers, e.g., the inference barrier between particulars and universals. Irrelevance-based counterexamples threaten these inference barriers: $Fa \models \forall x(Gx \rightarrow \exists y\exists z(Fy \wedge Gz))$ ([Russell and Restall 2010, p. 244](#)). The conclusion states that all G s are such that something is an F and something is a G . To illustrate that the *Universal Re-Gap Principle* avoids this example, let me sketch a state-based analysis. States that exactly exclude the conclusion are to the effect that there is a G but there are no F s. States that exactly exclude Fa are concerned with the individual a . Thus, no states that concern all F s exactly exclude the premise Fa .
- 27 [Fine \(2021, p. 899\)](#).
- 28 [Fine \(2021, p. 898\)](#).
- 29 The view that a moral fact N contains non-moral N -making facts is called *the wide view* of moral facts in [Cuneo and Shafer-Landau \(2014, p. 421\)](#).
- 30 I assume that for every two consistent states s and t that are incompatible s contains a state that exactly excludes t . Given this assumption, every counterexample falsifying the Gap Principle falsifies the Re-Gap Principle. Suppose descriptive D entails normative N while the descriptive components of at least one consistent state s incompatible with N is compatible with D . Since s is incompatible with D while s 's descriptive component is compatible with D , s is normative. By the assumption s contains an exact excluder of D . However, the exact excluder cannot be descriptive because s 's maximal descriptive part is compatible with D . Thus, a normative state exactly excludes D . Consequently, the Re-Gap is false.

6 The Modal Profile of Authoritative Normativity

A Relevant Restriction

Chapter 4 revealed that logical autonomy theses that are based on exception clauses bear an explanatory advantage. Satisfying exception clauses explain how spurious counterexamples to an autonomy thesis arise and why they are indeed spurious, i.e., compatible with the autonomy of normativity. In Chapter 5, I developed the Re-Gap Principle according to which descriptive premises do not entail a normative conclusion unless the conclusion's normative content is irrelevant to the entailment relation. Conclusions that meet the exception clause are normatively irrelevant. I have argued that Re-Gap Principle specifies an interesting, joint-carving sense of autonomy.

To scrutinize whether normativity is autonomous in the specified sense, one has to carefully examine *normative states* and their *modal profile*. It is important to acknowledge though that states are primarily a tool in our semantic framework. Accordingly, whether normative states are wholly relevant to the exclusion of descriptive states depends on what normative states are considered to be and how deep descriptive-normative entanglements run. In this section, I consider the option that conceptual distinctions in normative theorizing force us to acknowledge descriptive-normative entanglements that refute the universal claim that *no* normative state exactly excludes a descriptive state. My overall aim in this chapter is to defend a *restricted*, yet substantial Re-Gap Principle that holds even if certain descriptive-normative entanglements are presupposed. To that end, I substantiate the claim that conceptual distinctions in normative theorizing suggest deep descriptive-normative entanglements with respect to irreducibly thick and almost thin normative concepts. Recall that the classification of a normative concept as thick or thin depends on its descriptive application conditions. I then argue that restricting the Re-Gap Principle to what we may call thin normative propositions does not necessarily reveal anything distinctive about normativity. I conclude that an interesting

restricted Re-Gap Principle concerns normative propositions that are normatively particularly relevant.

Irreducibly Thick Normative Concepts

Familiar cases of descriptive-normative entanglements are thick normative concepts like MURDER or SELFISH that have substantial descriptive application conditions.¹ Suppose, for example, that Ann did not kill Sarah. Consequently, Ann did not murder Sarah. Suppose further that the concept MURDER can be analyzed into descriptive and normative components. The descriptive components plausibly involve the concept KILLING. The Re-Gap Principle tells us that the example does not refute normative autonomy iff the entailment relation holds just in virtue of the descriptive content. Given our assumptions, this is indeed plausible because the normative components of MURDER, e.g., base motives or wrongness, are perfectly compatible with the fact that Ann did not kill Sarah. In other words, they do not help to exclude the premise. Accordingly, the entailment relation holds just in virtue of the conclusion's descriptive content. Hence, we face a spurious counterexample.

It is worth acknowledging, however, that accounts that assume *irreducibly thick* normative concepts generate plausible candidates for *genuine* counterexamples.² It is beyond the scope of this book to go into the details of the debate about thick normative concepts. To get an impression of the concern, however, consider the thick normative concept TACTFUL. Suppose x is *tactful* if and only if x is T and is good *in virtue of* being T where ' T ' is a predicate ascribing a descriptive property.³ Consequently, [x is not T] entails [x is not tactful]. Yet, any state to the effect that x is tactful must be to the effect that x is good *in virtue of* being T which might be seen to be wholly relevant to the exclusion of a state to the effect that x is not T . The motivation for this view is that the in-virtue-of clause generates a tighter descriptive-normative entanglement than conjunction would do. In any case, it is not clear that this state is not wholly relevant. Accordingly, some views of *irreducibly thick* normative concepts might require normative states that exactly exclude descriptive states, which in turn suffices to raise doubts concerning the autonomy of normativity as a general claim.

Almost Thin Normative Concepts

The concept OUGHT is commonly assumed to be a thin normative concept. That is, it involves almost no descriptive application conditions. Meta-ethical distinctions between *perspectivism* and *non-perspectivism* about moral oughts, however, suggest that OUGHT can be subject to descriptive

truth-conditions as well.⁴ Roughly, perspectivist views hold that moral oughts are affected by epistemic circumstances while non-perspectivist views deny this. To get a better grip on the distinction, it might help to think of the perspectivist's OUGHT to be relative to the agent's epistemic setting, e.g., her evidence or knowledge, while the non-perspectivist's OUGHT might be thought of as ought from the God's eye perspective. It will be illuminating to consider a case where the views' consequences come apart. The following example in Jackson (1991) illustrates why the distinction matters for normative autonomy theses.

(Drug Case) Jill is a physician who has to decide on the correct treatment for her patient, John, who has a minor but not trivial skin complaint. She has three drugs to choose from: drug A, drug B, and drug C. Careful consideration of the literature has led her to the following opinions. Drug A is very likely to relieve the condition but will not completely cure it. One of drugs B and C will completely cure the skin condition; the other though will kill the patient, and there is no way that she can tell which of the two is the perfect cure and which is the killer drug. What should Jill do?⁵

Non-perspectivism about moral oughts entails that since either drug B or C completely cures the patient Jill ought to give that drug. Suppose drug C is the perfect cure. Thus, on the non-perspectivist's account, Jill ought to give drug C. By contrast, perspectivism entails that taking into account Jill's evidence (e.g., a 50 percent epistemic risk of killing the patient) she ought to give drug A. On the perspectivist's account, it is not the case that Jill ought to give drug C. This normative consequence is *entailed* by the fact that Jill lacks a particular epistemic stance toward the proposition that drug C cures the patient, e.g., she lacks sufficient evidence. Whether we face a genuine counterexample to an autonomy thesis depends on the details of the accounts, but suppose a lack of sufficient evidence is spelled out in terms of thresholds of assigned probabilities. The corresponding descriptive propositions entail the normative proposition that it is not the case that Jill ought to give drug C. Note that probability assignments might require some sort of rationality constraint so that the relevant propositions do not count as descriptive. To raise the outlined concerns, however, it suffices that there is some descriptive way of stating that Jill does not have any information about drug C. This suggests that even *prima facie* thin normative concepts can be seen as essentially descriptively constrained, so that some normative states at least could be seen to exactly exclude descriptive states.

One might ask whether the Re-Gap's exception clause excludes those cases because the perspectivist's OUGHT is a conjunction of epistemic conditions and an epistemically *unconstrained* OUGHT. Suppose that a perspectivist ought-claim just is a non-perspectivist ought-claim together with the fact that particular epistemic conditions are met. Since the unconstrained OUGHT is compatible with Jill's failure to meet the epistemic conditions, the normative content does not help to exclude the descriptive premise. Yet, this strategy only works if the perspectivist's OUGHT can be analyzed in terms of an unconstrained ought. Notice that in the scenario, there does not seem to be an epistemically unconstrained sense of OUGHT according to which Jill ought to give drug A. The perspectivist ought-claim is incompatible with the descriptive state that Jill does not meet particular epistemic conditions *because* that ought-claim is relative to Jill's epistemic situation. Thus, normative propositions that involve the non-perspectivist's OUGHT might turn out to be relevantly entailed by descriptive premises.

A Thin-Rather-Than-Thick Restriction and Its Limits

The examples considered so far illustrate that some normative concepts might turn out to be deeply entangled with their descriptive application conditions. As a consequence, one would face entailments where not only the conclusion's descriptive but also its normative content is relevant to the entailment relation. Entailments where the normative content is relevant, however, refute the Re-Gap Principle. To be clear, I do not intend to defend any particular view on whether (and if so which) normative concepts are entangled. (In fact, I'm skeptical about such entanglements.) The aim of this chapter is, however, to assess the consequences of such a view for the debate about the autonomy of normativity. I will argue that there are plausible ways of restricting the autonomy thesis so that we are still left with an interesting claim about the normative domain. This restricted principle will allow the proponents of such a view to defend an albeit weaker but relevant sense of normative autonomy.

One might wonder why I do not account for these examples by extending the exception clause of the Re-Gap Principle. It is important to acknowledge, however, that an exception clause is supposed to specify and explain conditions under which counterexamples are *spurious*. The considered examples, however, seem to be genuine. If they are successful, they show that the normative domain considered as a whole is not autonomous. That said, there might still be a normative subdomain that is autonomous in the sense specified by the Re-Gap Principle. So, at this point, I return to the strategy of providing a *restricted* autonomy thesis. That is, I consider restrictions of the domain to which the Re-Gap Principle applies. Note

that the more widespread descriptive-normative entanglements are, the less clear it is that any interesting normative subdomain is autonomous.

In what follows, I consider a strategy to restrict the Re-Gap Principle based on the distinction between thick and thin normative concepts. I argue that this strategy does not guarantee that one is still concerned with an interesting autonomy thesis. Instead, the Re-Gap Principle should be restricted to a normative subdomain that is normatively particularly relevant. The task of specifying such a domain and scrutinizing its essential traits is then undertaken in the remainder of this chapter.

Let me first scrutinize a restriction strategy based on the thick-thin distinction. How to capture a precise distinction between thick and thin normative concepts is controversial, but it is common to assume that for a normative concept to be *thin* it should involve little to no descriptive application conditions. We already encountered paradigmatic examples of thick normative concepts such as TACTFUL. On some accounts, however, even seemingly thin normative concepts turn out to be thicker than it is commonly assumed. I shall briefly present one example in order to motivate a restriction of the Re-Gap Principle to what I shall call thin normative propositions.

Consider the idea that not just everything counts as a moral system.⁶ For example, we might assume that as a matter of conceptual truth, every moral system must be such that genocide is morally wrong. Otherwise, it just does not count as a moral system; however, it might still count as a normative system, for example, a prudential system. Given this constraint, the descriptive premise [Act *a* is a genocide] entails the normative conclusion [*a* is morally wrong]. It is easily verified that under the assumption that this entailment relation holds, it does so in virtue of the normative content and thus refutes the Re-Gap Principle. I use the example merely for illustrative purposes. That is why I set aside other strategies to respond to it. For example, one might argue that genocide is a thick normative concept that is defined in normative terms or deny that there is a normative concept that applies to all instances of genocide as a matter of conceptual truth. Instead, I shall focus on strategies to restrict the Re-Gap Principle so that its domain of application excludes propositions like [*a* is morally wrong] where MORALLY WRONG is constrained in the outlined way. The crucial question is on which basis this can be done.

The outlined example suggests that one might draw on the distinction between thick and thin normative concepts. I shall say that a proposition is a thin normative proposition iff it is a normative proposition and involves only thin normative concepts.⁷ Restricting the Re-Gap Principle to thin normative propositions rules out examples like the genocide-example because, on the outlined account, MORALLY WRONG is a thick normative concept. On this strategy, however, the autonomy thesis is in danger of

becoming an overly weak claim. If thin normative concepts have no descriptive application conditions, thin normative propositions are not relevantly entailed by descriptive propositions *by assumption*.

One might object that while the Re-Gap Principle restricted to thin normative propositions is indeed analytically true, it is still a significant claim about normativity because thin normative concepts play a crucial role in normative theorizing.⁸ Note, however, that his reasoning requires:

Thin Relevance: Thin normative concepts are normatively particularly relevant.

Thin Relevance illustrates that to obtain a significant autonomy thesis, it is not enough to specify just any normative subdomain so that a restricted version of the Re-Gap Principle is true. It has to be shown that the normative subdomain that is autonomous in the specified sense is normatively particularly relevant (or at least somewhat interesting). This relevance-based strategy leaves it open whether possessing normative relevance crosscuts the thick-thin distinction. In the next section, I argue that a promising relevance-based strategy is to focus on authoritatively normative propositions.

Before turning to the notion of authoritative normativity, let me highlight that considerations along similar lines have been put forward against putative genuine counterexamples to the Is-Ought Gap.⁹ We already encountered Searle's argument that is supposed to show that we can infer obligations from valid promises. In response, it has been pointed out that promises merely suffice to generate obligations internal to a system of rules, e.g., promissory-rules. Any such internal obligation is arguably premised on the agent's accepting or having reason to accept the rule-system in question, whereas *substantial* obligations require that the system's normative significance is independent from the agent's accepting it. In the next section, I will discuss competing accounts how best to conceive of substantial obligations. So, let me restate the agenda of this chapter: specifying a normative subdomain that possess a particular form of normative relevance and scrutinizing conditions under which it is autonomous in the sense specified by the Re-Gap Principle.

What to Do: Authoritative Normativity

The *authority* of normative propositions or concepts recently gained growing attention in metanormative debates.¹⁰ Yet, whether there is a coherent concept of authoritative normativity and if there is, what exactly it tracks is controversial.¹¹ In this section, I introduce a preliminary notion of authoritative normativity and contrast it with alternative notions in the

debate. My preferred construal of authoritative normativity will have to wait until the discussion of how authoritative normativity bears on modal autonomy.

Authoritative normativity contrasts with *formal* normativity. While authoritative normativity is supposed to be inherently significant, formal normativity is the normativity displayed by *any* standard one can meet or fall short of like the rules of chess, or club rules.¹² For example, your failure to pray may fall short of religious standards. However, if you do not care about religion, this does not entail that your behavior is in any interesting sense defective, whereas falling short of a moral standard usually means that you are guilty of a mistake even if you do not care about morality. The view that there is a form of authoritative normativity is often contrasted with a view called *normative deflationary pluralism*. Basically, normative deflationary pluralism is the view that there are several normative systems and concepts that coexist, none of which is inherently more significant than the others. Eventually, I argue that deflationary pluralists either have to deny normative autonomy or have to confine themselves to weak autonomy theses, while those sympathetic to authoritative normativity can endorse a substantial autonomy thesis. However, this will have to wait until we have a better grip on authoritative normativity.

McPherson (2018) introduces the notion of authoritative normativity via the concept AUTHORITATIVE OUGHT. According to this construal of authoritativeness, normative propositions that exert authority are supposed to stop the process of practical deliberation by settling what-to-do-questions.¹³ McPherson presents the following conflict to motivate this characterization:¹⁴

(Sticky Situation) You find yourself in a sticky situation. You conclude that morality requires you to stay and help, while prudence dictates that you take the money and run. Torn, you ask yourself: given all of this, what ought I to do?¹⁵

Someone in Sticky Situation¹⁶ seems to ask herself an intelligible, non-trivial question. However, if one takes ‘ought’ to express a moral ought, or a prudential ought, then the question is trivial. The person in Sticky Situation takes herself to know what she morally (prudentially) ought to do. McPherson suggests interpreting the question asked in Sticky Situation in terms of what he calls *practical ought* which is supposed to be a *distinctively authoritatively normative* concept. Plunkett (2020) glosses the question at issue in Sticky Situation in terms of what we *really and truly should do*. Note that this proposal is compatible with a view that is called *moral rationalism*, according to which in cases of conflict, we always authoritatively ought to do what we morally ought to do. Indeed, McPherson

argues that we need authoritatively normative concepts to make sense of these kinds of views.¹⁷

Authoritative normativity is not limited to a particular kind of normative concept. We can consider concepts like AUTHORITATIVE REASON¹⁸, AUTHORITATIVE REQUIREMENT, AUTHORITATIVELY BETTER, etc. Note that these concepts do not play a settling role in normative thought, which is why it is a substantive commitment of views like the one in McPherson (2018) and McPherson and Plunkett (2024) that a settling-based characterization of authoritativeness can be used to provide a unified account of all authoritative normative concepts.¹⁹

To illustrate the difficulties in specifying a notion of authoritative normativity, let me go through some alternative characterizations that have been considered inadequate to capture the idea of authoritative normativity, or whose adequacy is controversial. First, one might think that authoritative normativity is captured by *categoricity*. If a norm is categorical, indifference to the norm does not evade its application. However, Foot (1972) pointed out that the rules of etiquette are categorical in this sense. It is impolite to interrupt someone even if you do not care. Since the rules of etiquette do not seem to possess normative authority, a characterization in terms of categoricity fails.²⁰ Second, according to Parfit (2011), there is a distinction between normativity in the ‘*reason-implying*’ and in the ‘*rule-implying*’ sense.²¹ Normativity in the reason-implying sense is Parfit’s way of adverting to authoritative normativity. However, it has been argued that we use notions like ‘legal reason’, ‘reason of etiquette’ or ‘religious reason’. Insofar as one thinks that these notions are not reasons in ‘the standard normative sense’,²² one seems to presuppose rather than explain the formal-authoritative distinction.²³ Third, and finally, one might think of authoritative normativity in terms of ‘*all-things-considered*’ characterizations. This proposal is particularly salient when one is trying to characterize the authoritative ought. McPherson rejects this proposal because he takes the ‘all-things-considered’-qualification to be read norm internal. Accordingly, one obtains claims about what we morally, prudentially, or by light of etiquette all-things-considered ought to do. If McPherson is correct, the ‘all-things-considered’-qualification does not solve conflicts like the one in Sticky Situation.²⁴

Howard and Laskowski (forthcoming) have recently argued that the demand for an additional authoritative ought, other than the all-things-considered ought, is based on a particular view of normative subdomains. They argue that an authoritative ought becomes necessary only if we assume that there are different normative standpoints or domains, such as morality or prudence, that determine how to balance the reasons in order to obtain the corresponding ought, for example, the moral ought. They contrast this view with what they call the *subset view*, according to which

the different normative flavors correspond to different subsets of the totality of reasons. In order to determine the moral ought, one must focus on the balance of the unique subset of moral-reason-giving considerations. According to the authors, on the subset view, the all-things-considered ought is uniquely comprehensive and thus settles the what-to-do question. At this point, I will remain neutral on whether ‘all-things-considered’-qualifications solve the problem. The aim of this chapter is to show that there is a normative domain that is particularly relevant and modally autonomous in the sense outlined in [Chapter 5](#).²⁵

Another contrast that might help to fix ideas concerns the strength of norms. Normative requirements can differ in their stringency. In [Hansson \(2014\)](#), strength is cashed out formally as follows: let O_1 and O_2 be normative operators or predicates. O_1 is stronger than O_2 iff for all P if $O_1(P)$ then $O_2(P)$, while the converse does not hold.²⁶ To see that Sticky Situation is not about strength thus defined, assume that you authoritatively ought to stay. If AUTHORITATIVE OUGHT is stronger than PRUDENTIAL OUGHT, then by definition, it would be prudentially required to stay. Yet, the whole point of Sticky Situation is that there is a norm conflict because prudence does not require you to stay.

The contrasts suggest a negative answer to the question of what authoritative normativity is. Concerning a positive answer, I will stay neutral with respect to the question of whether some authoritatively normative concept is primitive.²⁷ For the present purposes, it matters that someone who takes herself to know that she authoritatively ought to ϕ and wonders whether ϕ -ing is the thing to do does not count as competently using the concept.²⁸

A Non-Arbitrary Gap

In this section, I examine how the focus on authoritative normativity bears on the modal autonomy of the normative domain. I have argued that a restriction of the Re-Gap Principle to thin normative propositions yields the desired result that the restricted principle is not subject to counterexamples. Yet, by restricting the Re-Gap Principle to propositions that have no descriptive application conditions the autonomy thesis becomes an overly weak claim. In order to retain the Re-Gap’s substantiality, I shall examine whether authoritatively normative propositions are autonomous with respect to the descriptive domain. To that end, I consider two features of authoritative normativity discussed in the literature and explain why these features suggest that authoritatively normative concepts have no non-trivial descriptive application conditions.²⁹ The results of this section provide the ground for an interesting modal autonomy thesis that can be endorsed in case counterexamples to the unrestricted Re-Gap Principle were to succeed.

Extension-Revealing and Authority-Revealing Concepts

McPherson and Plunkett (2020) argue that there is a general tension between how *extension-revealing* and how *authority-revealing* normative concepts are.³⁰ In what follows, I introduce the distinction and argue that this tension is robust. The authors assume that a concept is extension-revealing to the extent that ordinary use of that concept tends to make facts about its extension accessible.³¹ Examples considered by the authors are among others BRAVE and POLITE. Extension-revealing concepts involve information about which particular entities, e.g., acts or persons, fall under that concept. The more extension-revealing the more substance-orientated the concept is. For example, the concept BRAVE plausibly applies only to acts performed at (perceived) risk. On the other hand, a concept is authority-revealing to the extent that competence with that concept tends to make its authority accessible.³² Paradigm examples are AUTHORITATIVE OUGHT or AUTHORITATIVE REASON. At first glance, one might think that MUST is more authority-revealing than SHOULD, for MUST is more binding than SHOULD. However, recall that authority is not the same as stringency. $[x \text{ must } \phi]$ entails $[x \text{ should } \phi]$ but not vice versa. Thus, MUST is stronger than SHOULD. By contrast, authority matters in cases of (perceived) symmetry. Suppose that somebody's dress looks awful. You think that it would be tactful to lie but sincere to tell the truth. If one of the concepts TACTFUL and SINCERE is more authority-revealing, then the corresponding judgment helps to settle the what-to-do-question.³³

Both features allow for a merely epistemic and an objective reading. On the epistemic reading, the question is how transparent the extension or the authority of particular concepts is to us; on the objective reading, the question is whether the concept objectively determines its extension in a substance-orientated way or whether the concept objectively exerts authority in cases of perceived normative conflict. That is, conceptual competence with the concepts requires to acknowledge substance-constraints or authority orderings, respectively. Since the autonomy debate concerns entailment relations, I set aside cases where concepts are extension- or authority-revealing because of the way people have been trained to use the concept over time. Instead, I shall focus on *conceptual truths* that reveal a concept's extension or authority. The crucial question is whether there is a robust tension between the two features thus understood.³⁴

Arguments in support of this tension bear some resemblance to Moorean open-question arguments.³⁵ Note that I do not take arguments of this kind to establish the conclusion that authoritatively normative concepts cannot be extension-revealing. Instead, I take them to show that it is unclear how descriptively constrained concepts should establish their authority. After presenting one variant of the argument, I turn to another argument

that aims to show that authority is indeed incompatible with descriptive application conditions. The second argument is based on substantial assumptions about the nature of authoritativeness.

Consider an arbitrary authoritative normative concept N such that accepting that ϕ -ing is N while wondering whether ϕ -ing is the thing to do involves a conceptual mistake. Suppose further that N is also extension-revealing. That is, N might have substance-orientated necessary or sufficient conditions. Thus, N involves information to which substances it applies. Suppose that N applies to all entities that are D where D is some descriptive condition. Similar reasoning applies to descriptive necessary conditions. A first thing to clarify is how D relates to N 's authority. I can think of two possibilities: (i) N exerts authority *and* applies whenever an act is D , call this the *combinational variant* or (ii) N exerts authority and it does so partially in virtue of applying to acts that are D , call this the *integral variant*.

At this juncture, it is important to note that both variants seem to be in conflict with previous assumptions. First, consider the *combinational variant*. Since N is assumed to be authoritative, N -truths settle what-to-do-questions by means of tie-breaking asymmetries. Suppose, for concreteness, that TACTFUL is such an authoritative normative concept. Let the descriptive condition D be such that acts by which *the agent shows consideration for other's feelings* are tactful. If D is merely a conjunctive condition, we can consider an alternative normative concept N' —let's say SHMACTFUL—that exerts authority without applying to all D -acts. Accordingly, acts by which the agent does not show consideration for other's feelings can be shmactful. As far as authority is concerned, N and N' , in this case TACTFUL and SHMACTFUL, are on a par. This symmetry is incompatible with the assumption that N -truths settle what-to-do-questions. To see why, consider two alternative courses of action. The agent shows consideration of other's feelings only by one of them. This course of action is tactful, but the other might still be shmactful. Since both concepts share their normative components, it is unclear how the concept TACTFUL should invoke the required asymmetry. By contrast, the *integral variant* seems to have the resources to respond to that issue. If N exerts authority partially in virtue of applying to acts that are D then the alternative concept N' seems to miss an important trait of authority which in turn breaks the symmetry between N and N' . If the concept TACTFUL is authoritative partially in virtue of applying to all acts by which the agent shows considerations for other's feelings, then we miss something about the answer to the what-to-do-question when we pursue the shmactful option. This, however, suggests that D exerts authority. What to do is (at least partially) settled by the fact that one of the acts is an act by which the agent shows considerations for other's feelings. We thereby locate authority in the wrong place. Authority, if it is coherent, is an essential trait

of normativity. Consequently, if *D* exerts authority, it turns out to be normative contrary to our assumptions.³⁶ Hence, we either locate authority in the wrong place or unpleasant symmetries between alternative normative concepts like TACTFUL and SHMACTFUL keep authority questions open.

Non-Arbitrariness

Let me now turn to an argument that concludes that the authority of a normative concept is incompatible with non-trivial descriptive constraints. The argument assumes that authority is incompatible with arbitrariness. In a nutshell, the thought is that any choice of descriptive application conditions is subject to a form of arbitrariness that is incompatible with authority. Before I present the argument in detail, let me clarify the notion of arbitrariness and the assumption that authoritative normativity is essentially non-arbitrary.³⁷

To illustrate why non-arbitrariness plays an important role in the context of authoritative normativity, reconsider Sticky Situation. Assume that either prudence or morality is authoritative in this situation. McPherson points out that flipping a coin to decide between prudential and moral ought does not bring us any closer to what we authoritatively ought to do.³⁸ Similarly, following the rule of always acting as morality requires involves an arbitrary rule-selection given the symmetry between this rule and the corresponding prudential rule of always acting as prudence requires. That prudence or morality favor the prudential or the moral rule, respectively, does not help neither given the assumed symmetry between prudence and morality. In order to tell which obligation is authoritative, we have to reach some point at which an objective, stance-independent, non-arbitrary asymmetry between the two systems occurs. The arbitrariness that is involved in prioritizing one over the other without an explanatory asymmetry between them seems to be incompatible with the claim that some ought exerts objective, stance-independent authority.

A non-arbitrary distinction between, for example, morality and prudence requires an underlying asymmetry. This requirement, however, seems to invoke an infinite regress. One way to stop the regress is by appeal to the nature of one of the systems. If one of them is by its nature and thus essentially asymmetry-introducing, this would plausibly cancel the requirement of more fundamental asymmetries.³⁹ Suppose that there is a normative system or concept that is essentially asymmetry-introducing. The claim is that the non-arbitrariness assumption rules out that the relevant concept or system is non-trivially descriptively constrained. The reason is that any non-trivial descriptive constraint works against the non-arbitrariness of the concept. No constraint can recommend itself. Thus, it re-opens the regress. Any choice of constraints would again be arbitrary.

One might think that the following line of argument works: the descriptive constraint is not arbitrary because it matters, for example, to the flourishing of human lives. However, this line of reasoning either presupposes a normative premise (e.g., that the flourishing of human lives is valuable) based on a normative concept that is arbitrarily chosen or it assumes that normative authority is not mind-independent in the way that objectivists usually assume.⁴⁰

Before turning to objections against authoritatively normative concepts, let me reconsider deflationary normative pluralism. According to this view, there is no ultimate answer to what-to-do-questions that does not involve non-arbitrary selection at some point. Deflationary pluralism can also be found in non-normative debates, for example, in the debate about mathematical concepts such as the concept of a set. While there might be aesthetic or pragmatic reasons to favor a particular concept, deflationary mathematical pluralists assume that there is no objectively privileged concept of a set.⁴¹ What distinguishes the mathematical and the normative debate, however, is that in the mathematical debate, there is no analogue to the pressing what-to-do-question. Arbitrariness does not pose an analogous threat to deflationary mathematical pluralism. Notice that the normative deflationary pluralist is free to assume constrained as well as unconstrained normative concepts, but there is no objective, stance-independent way for her to assign a privileged status to either of them. Thus, some normative subdomains might be autonomous, others might fail to be, but this conjuncture does not reveal anything of particular interest about normativity.⁴²

Let me take stock. If the arguments are correct, then authoritatively normative concepts are not non-trivially descriptively constrained. This makes an autonomy thesis that is restricted to authoritative normativity promising. It is important, however, what motivates the restriction. It matters to the substantiality of the resulting autonomy thesis that there are independent reasons to think that what is particularly relevant to normative deliberation is what we authoritatively ought to do. Before turning to the autonomy debate, I discuss some objections raised against authoritative normativity.

Work for a Theory of Authoritative Normativity

In this section, I turn to objections against the coherence and theoretical utility of authoritatively normative concepts. My discussion is based on objections raised by Baker (2017, 2018). It would go beyond the scope of this book to offer a comprehensive defense of authoritative normativity. However, Baker flags that his criticism is ‘part of the venerable tradition of trying to get philosophers to stop being so easy on themselves’.⁴³ It is this more modest aim I hope to make progress with: authoritative normativity

does not require rashness. It is worth emphasizing that the aim of the relevance-based approach of authoritative normativity that I am going to propose is to elucidate the concept in terms of orderings and corresponding hierarchies that are familiar tools in non-normative debates. This project is distinct from the project of defending the claim that authoritative normativity is not non-trivially descriptively constrained. I pursued the latter project in the previous section. So, one need not adopt the relevance-based approach to endorse the autonomy of the authoritatively normative domain.

Further Suspect Objection

One objection to authoritatively normative concepts such as AUTHORITATIVE OUGHT is that they cannot solve the perceived conflict in cases like Sticky Situation.⁴⁴ Baker claims that the conflict results from ‘too many oughts, not too few’.⁴⁵ Adding further oughts, so the thought goes, if anything, makes the situation worse. The supporting argument is the following: the initial symmetry between what we *morally* ought to do and what we *prudentially* ought to do cannot be adjudicated by introducing a further ought because it would be subject to the same symmetries. In response to this objection, I shall do two things. First, I suggest an interpretation of AUTHORITATIVE OUGHT that invites the further-suspect-objection and explain why it is misguided. Second, I propose an alternative interpretation and argue that it is not subject to the objection.

One might think that an authoritative ought is an accumulated ought in the sense that it results from the accumulation of different standards.⁴⁶ For concreteness, recall the oughts involved in Sticky Situation. According to prudence, you ought to take the money and run; according to morality, you ought to help. Suppose that the authoritative ought is the result of a weighted accumulation of these standards. In the example, morality might be weightier than prudence and thus you authoritatively ought to help. This, however, requires an accumulation function. In other words, accumulation requires a standard that guides the accumulation. AUTHORITATIVE OUGHT thus understood, indeed raises the question of how the accumulation function can be assumed to settle what-to-do-questions. We took the conflict in Sticky Situation to arise in virtue of a perceived symmetry between the standards of prudence and morality. AUTHORITATIVE OUGHT now appears to be the result of *another standard* (the accumulation function). This standard is in turn plausibly subject to symmetries concerning alternative accumulation functions. In particular, the moral (prudential) ought might be itself the result of an accumulation function that gives dominant weight to the moral (prudential) standard. We indeed face a further suspect.

To see why this interpretation is misguided, consider the following analogy based on a non-metaphorical usage of the term ‘authority’. Suppose you are a soldier. You find yourself in a risky situation and wonder what to do. Suppose further three other soldiers are around you. Two of them command to attend, while the third commands to advance. One option in order to settle what to do is to accumulate the commands. By putting equal weight on each of them you obtain the result that you ought to attend. This strategy requires a standard of accumulation, e.g., putting equal weight on each of the commands. The accumulated result might indeed be seen as a forth ought or a forth standpoint. By the standards of each of the soldiers, you either ought to attend or advance, respectively; by some standard of accumulation, you ought to attend. Now suppose that one of the soldiers has a superior military rank. The soldier that commands to advance is a major. In that case, it is inappropriate to accumulate the commands. What matters to the what-to-do-question is that one of the commands trumps all other commands. The major’s command is *prior* to the other commands in virtue of an *internal feature* of the command’s source.

I propose that authoritative normativity is best understood in terms of normative relevance orderings. This is important to do justice to the symmetry-breaking nature of AUTHORITATIVE OUGHT. In Sticky Situation, one wants to know whether morality or prudence is normatively more relevant and in particular, what is normatively most relevant in order to decide what to do. We authoritatively ought to ϕ if the ought that is normatively most relevant requires to ϕ . Accumulated oughts might very well be part of such an ordering. However, to establish what we authoritatively ought to do is to establish which ought is the maximum with respect to a normative relevance ordering. It is this interpretation of AUTHORITATIVE OUGHT that accounts for the asymmetry-introducing nature of authoritatively normative concepts. Suppose that we authoritatively ought to do what morality requires. The accumulation-interpretation must assume an accumulation function that puts dominant weight on a moral ought. Thus, there is indeed a further ought in addition to the moral ought: the ought of a standard introduced by the accumulation function. By contrast, the relevance-interpretation requires that in the described scenario moral oughts are superior to all other oughts because moral oughts possess the highest rank in the normative relevance ordering.⁴⁷

Unmotivated Objection

Another objection concerns the question of what exactly is shown by examples like Sticky Situation. Proponents of authoritatively normative concepts argue that the question asked in Sticky Situation should be interpreted in terms of an authoritative ought. By contrast, Baker argues that a

skeptic can say that the question should be interpreted in terms of motivation. Someone who asked what she ought to do in a situation where she already accepts conflicting oughts concerning different standards is asking for an ought-claim that will *move* her to act.⁴⁸ So, in cases like Sticky Situation, the agent just fails to be moved by morality or prudence and seeks for something that will move her. To see that this move will not solve the issue consider a modification of Sticky Situation.

(Post Sticky Situation) You found yourself in a sticky situation. You concluded that morality requires you to stay and help, while prudence dictated that you take the money and run. You decided to act according to your moral obligations. So, you stayed and helped. Later you ask yourself: given all of this, what was I obliged to do?

In Post Sticky Situation the person already knows which ought-claim moved her (the moral one). What she wants to know is whether she was moved by the right, that is, the authoritative ought-claim. As far as I can see, there is no plausible interpretation of the question asked in Post Sticky Situation in terms of motivation. What underwrites the motivation-based reading is the claim that an authoritative ought is ‘unnecessary, so long as we are constituted to care about some oughts more than others’.⁴⁹ However, one must keep apart two importantly different questions. One is about the motivation of agents and how accepting ought-claims relates to their motivation to act accordingly; the other one concerns whether what-to-do-questions can be objectively settled. Even if someone agrees that she really and truly should help, cases of *acrasia* where agents act against their better judgment are still possible.

Connections and Theoretical Role

The last objection that I will discuss here states that authoritatively normative concepts are neither linked to more familiar theoretical notions nor are they theoretically useful. Part of the objection is that reference to ‘authority’ is metaphorical, vague, and unilluminating. To respond to this objection, I shall do three things. First, I show that our understanding of the concept of authority in non-metaphorical cases informs our understanding of concepts like *AUTHORITATIVE OUGHT*. As will become clear, the non-metaphorical usage helps to clarify the role of authoritatively normative concepts, for example, with respect to normative criticism. Second, I argue that on a relevance-based approach, authoritative normativity can be understood in terms of the familiar concept of relevance and corresponding orderings. On this account, normative authority relates to a

theoretical notion that is familiar from debates that are commonly not taken to be normative. Third, I defend the claim that the notion of authoritative normativity structures views on normativity.

Let me start with the claim that authority is objectionably obscure. Consider a paradigm case of authority-related issues:⁵⁰ the sergeant commands to attend while the major commands to advance. Since the major possess authority, the soldier should obey the major's command. This situation does not seem to be puzzling or obscure. Moreover, we can describe the situation as follows: usually the soldier should obey the sergeant's commands. If she refrains, then she is criticizable. However, if there is a superior military rank, such as the major, that commands otherwise, then she should obey the major's command. Again, if she refrains, then she is criticizable. Yet, she is no longer criticizable for not obeying the sergeant's commands (in cases of conflict). Note that the major's command did not annihilate the sergeant's command, but it did annihilate a source of criticism.

One way of thinking about AUTHORITATIVE OUGHT is that it annihilates normative sources of criticism, while leaving the relative oughts in place. So, if you authoritatively ought to help, then you still prudentially ought to run, but your non-compliance with the prudential ought is not criticizable. One might object that even if one authoritatively ought to help, one is still criticizable from the prudential perspective. To discuss these objections in detail would require more space than I can spend here, but to me, it seems that somebody who acts as she really and truly should do is an important sense void of critique. This is not to say that we need to understand authoritatively normative concepts this way. Yet, the authority-metaphor seems to contribute to our understanding of the phenomenon at issue.

If analogies like the previous one hold, then metaphorical characterizations in terms of *authoritative*, *trumping*, and *overriding* oughts help to elucidate authoritatively normative concepts. However, it seems fair to say that these characterizations are not sufficient for a compelling account of authoritative normativity. Let me quote a passage in Baker (2018, p. 579, italics in original):

Let's say morality is authoritative and etiquette is not. Let's say further that in this case they conflict, the first telling me to confront my host's racist bile, and the second to change the subject. It is natural to think that the authoritative prescription *overrides* or *trumps* the merely formal prescription. But what does this mean? Given our stipulation, it does not mean that the action is not impolite.

First, note that the previous analogy suggests a response: it is still impolite to confront the host, but one is no longer criticizable for doing so.

Given the stipulation, the moral ought is the *relevant one* for decision-making; moral oughts are more relevant to the what-to-do-question than the oughts of etiquette. Insofar as the moral ought is the authoritative one, it is even most relevant.

While it might seem that locutions like ‘decisiveness’ or ‘significance’ are equally or even better suited to capture authority, I will argue that an account in terms of normative relevance is a theoretically fruitful way to approach normative authority. The notion of relevance plays a theoretical role in many debates including causation, explanation, and theory of language. One might insist that it is not clear that the debates refer to the same notion of relevance. However, a point that I want to stress is that the debates provide us with a grip on the notion of relevance that is not normatively biased in the way that the notion of decisiveness might be. A relevance-based account is thus less prone to circularity objections. In particular, the distinction between a *pragmatic* and an *objective* notion of relevance is discussed in the mentioned debates.⁵¹ As will emerge soon, this distinction matters for the perspectives on normativity that are at issue. The account that I am going to propose is based on an ordering of normative propositions with respect to normative relevance. Based on this relevance ordering, authoritative normativity can be explicated in terms of what is normatively most relevant. This strategy, however, is less intuitive when framed in terms of significance or decisiveness. The claim that normative propositions possess lower or higher degrees of significance or decisiveness is less appealing. One might also wonder whether we should think of authoritative normativity in terms of *practical relevance*, for authoritatively normative truths tell us what to *do*. Yet, restricting normative relevance to practical relevance, I think, narrows down the applicability of authoritatively normative concepts. In general, there is no reason to exclude, for example, authoritatively normative concepts that settle what-to-believe-questions.

I already touched upon another strength of the normative relevance-approach in the context of the further-suspect-objection. One might think that ranking the different oughts requires an additional system or standard of normative relevance that allows us to evaluate the oughts. However, on the normative-relevance approach, normative propositions come with a special feature possessed in different degrees, namely objective normative relevance.⁵² This is a substantial assumption about normativity, but this is what makes the resulting autonomy thesis eventually a substantial and interesting claim. Given that normativity is such that normative truths possess objective normative relevance, authoritative normativity does not introduce further standards. Authoritative normativity represents what is normatively most relevant.⁵³

To address the third objection concerning the theoretical utility, let me first explain why Baker thinks that there is no theoretical work for AUTHORITATIVE OUGHT. Baker claims that in cases like Sticky Situation there is no *theoretical* problem that AUTHORITATIVE OUGHT helps to characterize. Morally the person ought to help, while prudentially the person ought to run. To put it in Baker's terms: 'This is how the world is. There is no mystery here crying out for explanation'.⁵⁴ He thereby indirectly characterizes a view on normativity. Roughly, different normative systems impose different, sometimes conflicting obligations on agents, yet there is no objective way to ultimately settle what-to-do-questions. This is *deflationary pluralism*. Yet, we might disagree: we might think that, while this is how the world is, this is not the whole story (not just for obvious reasons). A complete normative description of the world not only reveals what we morally and prudentially ought to do, but also which of them (if any) is normatively most relevant.

The outlined disagreement generates a job description for authoritatively normative concepts: they help to characterize more precisely the difference between two competitive views on normativity. The crucial question is whether normativity settles what is objectively prior concerning normative relevance. Baker's skeptics assume that 'we face a host of formally normative systems each with its own different form of psychological or sociological *relevance*'.⁵⁵ Thus, the skeptic belongs to the camp that answers this question negatively, while her competitors answer it positively. We thereby get at the heart of the disagreement between the two camps: one camp appeals to *(inter-)subjective relevance*; the other camp endorses a form of *objective normative relevance*.

This concludes my defense of the theoretical utility of authoritatively normative concepts. While my arguments do not provide a comprehensive defense of authoritative normativity, I hope to have shown that there is work for authoritative normativity and that normative relevance orderings play a crucial role in characterizing this work. Next, I investigate the modal profile of authoritative normativity and a restricted Re-Gap Principle.

Normative Authority and Its Modal Profile

At the end of [Chapter 5](#), I have shown that the Re-Gap Principle is true iff no normative state exactly excludes a descriptive state. Previous considerations in this chapter revealed that some normative concepts might be descriptively constrained so that the corresponding normative states exactly exclude descriptive states. In this section, I shall scrutinize normative states that make authoritatively normative propositions true. In particular, I shall investigate their mereological and modal profile. It would go beyond the

scope of this book to offer a semantic account of authoritatively normative propositions. However, I hope to make progress with a better understanding of some of their formal features.

Authoritatively Normative States

In the last section, I proposed to conceive of the claim that we authoritatively ought to ϕ as the claim that the normatively most relevant ought (in a particular situation concerning a particular agent) requires to ϕ . This interpretation assumes an objective partial pre-ordering of normative propositions with respect to their normative relevance.⁵⁶ I shall stay as neutral as possible with respect to the semantics of authoritatively normative propositions.⁵⁷ However, some assumptions will be required in order to do justice to the essential features of authoritative normativity discussed so far.

To account for states that make authoritatively normative propositions true, I shall use the following notational conventions. We assume that there is a subset of the set of normative states that comprises all authoritatively normative states. $n!$ denotes an authoritatively normative state. n_i denotes a normative state relative to a normative system i . The set of all possible normative systems is I , OS denotes the set of options of a particular agent in a particular situation. Finally, \preceq denotes a normative relevance pre-order on the set of normative states.⁵⁸

Let me briefly pause to clarify the methodology. It is important to acknowledge that the formal tools of a state space are neutral with respect to substantial questions about authoritative normativity. Likewise, conclusions drawn within the formal framework are always in need of plausible interpretations at the substantial level. What I will do is to consider different assumptions at the substantial level, examine how they might be accounted for in the state space, and show that the formal account helps to make theoretical options and assumptions they will require precise. Eventually, it will emerge that even the deflationary pluralist can adopt the state-based account to make the contrast between her position and proponents of authoritative normativity formally precise.

My first aim is to clarify what it takes for an authoritatively normative proposition to be true. That is, what it takes for $n!$ to obtain. To that end, let me go through Sticky Situation. Suppose OS comprises all options for a particular agent A . I will assume that to help and to run exhaust the options. For simplicity's sake, I assume further that there are no other oughts beside what A morally ought to do o_M and what A prudentially ought to do o_P . Suppose that A authoritatively ought to help and that we do not face a normative dilemma. According to the relevance-based

account, this is the case if the normatively most relevant ought prescribes helping for A . For that to be the case, the moral ought must be normatively more relevant than the prudential ought. Formally, this is the case iff $o_P \preceq o_M$ and $\neg(o_M \preceq o_P)$. This motivates the following truth-conditions for authoritatively normative propositions. Let w denote the actual world. A authoritatively ought to ϕ iff with respect to all *actual* oughts concerning A the normatively most relevant prescribes ϕ -ing for A , where ϕ -ing is a fully determined course of action.⁵⁹ Formally, A authoritatively ought to ϕ iff

$$o_i^\phi \sqsubseteq w \text{ iff } \max_{\preceq} \{o_i^\phi : \phi \in OS, i \in I\} \text{ prescribes } \phi\text{-ing.}$$

Let me now turn to n 's mereological profile. To assess whether a restricted Re-Gap Principle holds, the *modal* profile of authoritatively normative states must be scrutinized. Yet, I already showed that the mereological profile of states bears on their modal profile. So, two questions must be addressed: do authoritatively normative states contain (i) non-authoritatively normative states (henceforth: normative states) and (ii) descriptive states?

Let me start with considerations at the substantial level. On the face of it, it might seem plausible that if moral rationalism is true, propositions about what we authoritatively ought to do analytically contain propositions about what we morally ought to do. It has been argued, however, that doubts concerning the normative significance of morality are not *conceptually* confused. Accordingly, [Ann morally ought to ϕ] is not part of the content of [Ann authoritatively ought to ϕ]. A formal account of authoritative normativity should be able to account for this assumption. So, the mereological profile of authoritatively normative states must be constrained accordingly. For all normative states n_i : $n! \not\sqsubseteq n_i$. It is important to acknowledge though that authoritatively normative and normative states might necessarily co-obtain. If moral rationalism is true, then moral truths and authoritatively normative truths are intensionally equivalent. Again, this highlights the importance of hyperintensional distinctions between propositions in normative debates.

Second, I shall address the question of whether authoritatively normative states contain descriptive parts. As before, I start with considerations at the substantial level. If my arguments in the previous sections are correct, authoritatively normative concepts do not have *non-trivial* descriptive application conditions. It is now time to say a bit more about this qualification. Consider the proposition [Sarah authoritatively ought to help]. Part of the content of this proposition is arguably the descriptive proposition [Sarah exists]. Let me explain why I think that the existence-constraint is merely trivial. For one thing, it does not undermine non-arbitrariness-constraints. The what-to-do question that is settled by the

authoritatively normative proposition is already relative to a particular agent, in this case Sarah. For another thing, the proposition [Sarah likes ice-cream] presupposes her existence in the same way. The constraint is thus a purely formal constraint.

To account for trivial descriptive application conditions in the state space, I shall introduce the notion of a *tied* descriptive part.⁶⁰ A similar idea guided the test I used to argue that normative states do not exactly exclude the agent's non-existence. Consider normative states that are opposed to each other, in the way that, for example, the state that Sarah ought to stay is opposed to the state that Sarah is permitted to go. Earlier, I described normative states that are opposed to each other in that way as making opposing normative contributions. The two opposed normative states both contain the state that Sarah exists. It is their normative status that is responsible for their incompatibility. This constitutes an intuitive sense in which a descriptive part of a normative state is *tied* as far as normativity is concerned.⁶¹ Within the truthmaker framework we might model the notion of a tied part *d* of a consistent normative state *n* by requiring that *d* is part of all consistent normative states that exactly exclude *n*. With this notion of a tied descriptive part at hand, I introduce a second constraint on the mereological profile of authoritatively normative states: authoritatively normative states contain only tied descriptive parts.⁶²

The Modal Profile of Authoritatively Normative States

Let me now turn to the modal profile of authoritatively normative states. In order to assess a restricted Re-Gap Principle, I have to examine whether authoritatively normative states exactly exclude descriptive states. To that end, I shall scrutinize modal relations between authoritatively normative states and (i) normative states and (ii) descriptive states.

I start with the substantial level. Suppose moral rationalism is true. Accordingly, [A morally ought to ϕ] entails [A authoritatively ought to ϕ]. Suppose further that some version of a utilitarian theory is true so that [A's ϕ -ing maximizes happiness] entails [A morally ought to ϕ]. Given these assumptions and by transitivity, [A's ϕ -ing maximizes happiness] entails [A authoritatively ought to ϕ]. For the restricted Re-Gap Principle to hold, these cases must be ruled out. In what follows, I will discuss two ways of specifying the modal profile of authoritatively normative states that allow for a restricted Re-Gap Principle.

One option is to deny views like moral rationalism. On the level of states this means that the normative relevance pre-ordering is contingent.⁶³ On this view, in some possible worlds prudence might trump morality, while in others morality trumps prudence. To see why this allows to hold on to the autonomy of the authoritatively normative domain, consider

descriptive states d_1, d_2 that necessitate normative states n_P^ϕ, n_M^ψ , respectively. Neither of d_1 and d_2 necessitates n^ϕ because neither the descriptive states nor the normative states fix the normative relevance ordering. That is, they do not fix that either $n_P^\phi \preceq n_M^\psi$ or $n_M^\psi \preceq n_P^\phi$. Given this modal profile, authoritatively normative states are modally independent from descriptive and normative states except for incompatibilities due to tied descriptive parts. Thus, authoritatively normative states do not exclude descriptive or normative states. However, one might take issue with the consequence that this option requires a commitment to metaphysically *contingent* facts about normative relevance. I admit that it is a theoretical cost of this view that it requires that normativity endows each possible world with a normative relevance ordering.

Let me consider another way of making sure that the considered example is not valid. One can allow for views like moral rationalism, if it is assumed that normatively most relevant oughts are not descriptively constrained. And this is indeed plausible. If, for example, moral oughts are necessarily what is normatively most relevant, then moral oughts rule out that there are oughts that are ranked higher in the normative relevance ordering. Thus, in all possible situations obtaining moral states trump all other obtaining normative states. If my arguments are correct, then authority is incompatible with non-trivial descriptive application conditions. This suggests that normative states that are ranked high in the relevance ordering contain only tied descriptive parts. Accordingly, there is a modal gap between the set of authoritative and highly ranked normative states on the one hand and the set of descriptive and lowly ranked normative states on the other hand.

There is, I think, one great advantage of the proposed account of authoritative normativity. Recall that authoritative normativity helps to characterize competing views on normativity. On the one hand, there are those views that take normativity to objectively settle what-to-do-questions. On the other hand, there are views like normative deflationary pluralism that assume that no answer to that question goes beyond psychological or social relevance. However, even those views might want to admit that the question asked in Sticky Situation is intelligible. This leaves them with two options. Either those views adopt a form of error theory according to which both propositions [A authoritatively ought to help] and [A authoritatively ought to run] are false or they take both propositions to be true so that one faces a dilemma. There is a natural way within the state space to accommodate each of these views. That A authoritatively ought to ϕ is defined in terms of maxima with respect to \preceq . If all actual oughts are on a par, then the conflict transmits to the authoritative level because several oughts are most relevant. This models the dilemma option. Note that this allows the deflationary pluralist to admit that in no-conflict situations,

one authoritatively ought to act as the actual oughts prescribe. However, one can also define AUTHORITATIVE OUGHT in terms of a maximum with respect to a strict order. Let $s < t := (s \preceq t \wedge \neg(t \preceq s))$. If deflationary pluralism is true, there is no maximum with respect to $<$ and thus all authoritatively normative propositions are false.

One might object that given that normative states ranked high in the normative relevance ordering contain only tied descriptive parts, we could have restricted the Re-Gap Principle to, e.g., moral propositions right away. As has been emphasized several times during this chapter, however, it is not at all difficult to specify a domain so that propositions that belong to that domain are not entailed by descriptive propositions. One just has to collect all propositions that are not entailed by descriptive propositions. The crucial question is whether there is something distinctive about normativity so that *it* makes it the case that these normative propositions are not entailed by descriptive propositions. Just stipulating that morality is not descriptively constrained is not playing fair with respect to opponents of normative autonomy. However, if the previous arguments are correct, authoritative normativity is not only crucial for normative theorizing but one of its characteristic traits, i.e., that authority is non-arbitrary, guarantees that authoritatively normative propositions are not entailed by descriptive propositions. The mereological and modal profile of authoritatively normative states as presented here suggest that a restriction of the Re-Gap Principle to authoritative normativity is an interesting modal thesis about normativity that one can hold on to in case the general principle would turn out to be false.

Restricted Re-Gap Principle: Descriptive propositions do not entail authoritatively normatively relevant conclusions.

Summary and Prospects

This concludes the second part of this book. Explications of the logical autonomy thesis that exclude spurious counterexamples often fail to specify a substantial sense of autonomy. In some cases, the specified sense of autonomy is not substantial because the autonomy thesis becomes an overly weak claim. For example, conservativeness explications specify a sense of autonomy that does not reveal anything distinctive about normativity (see [Chapter 4](#)). According to them, all kinds of domains are autonomous with respect to their complements. In other cases, the specified sense of autonomy is more ambitious, but does not entail satisfying distinctions between spurious and genuine counterexamples (see [Chapter 5](#)). Transitions that, if we grant their validity, intuitively refute the logical autonomy

of normativity are compatible with those explications. For example, the Finean Gap Principle faces the objection that his account classifies intuitively genuine counterexamples as spurious. My first aim had been to develop an explication of the logical autonomy thesis that specifies a sense of autonomy that is not subject to spurious counterexamples, yet sensitive to putative genuine cases. The Re-Gap Principle meets these requirements.

- No descriptive premises entail a normatively relevant conclusion.

I have argued that the Re-Gap Principle captures a substantial sense of logical autonomy. The Re-Gap Principle is ambitious in requiring that no normative state exactly excludes a descriptive state. However, distinctions in normative theorizing suggest that some normative concepts might involve descriptive-normative entanglements that require normative states that are wholly relevant to the exclusion of descriptive states. Yet, this does not mean that there is no restriction of the Re-Gap Principle that does justice to the pre-theoretical idea of normative autonomy. I have argued that we should focus on authoritative normativity. If my arguments are correct, authoritative normativity is incompatible with arbitrariness and thus merely trivially descriptively constrained. I proposed a way to approach authoritative normativity in terms of normative relevance orderings. The discussion revealed that one can hold on to a substantial sense in which normativity is logically autonomous in case the more general thesis turns out to be false. The restricted Re-Gap Principle focuses on two dimensions of normative relevance, namely normatively relevant entailments and normative relevance orderings.

- Descriptive premises do not entail authoritatively normatively relevant conclusions.

Proponents of normative autonomy might hesitate to adopt that autonomy thesis. After all, it is compatible with there being valid inferences from descriptive premises to a normative conclusion. I hope to have shown in this chapter that this might be a consequence of a *modal* explication of an autonomy thesis that we must swallow, if we are not merely interested in weak senses of autonomy. To be clear, I think that the Re-Gap Principle is a promising candidate, but if it turns out that some normative distinctions require distinctively descriptive constraints this strikes me as normatively relevant. We would have to face cases where descriptive propositions entail normative propositions in virtue of the normative content. Fortunately, this is not the end of this book. In its last part, I turn to metaphysical autonomy theses. I will argue that the most joint-carving autonomy thesis is an *explanatory* thesis: roughly, no normative proposition is true *in virtue of* descriptive propositions.⁶⁴ Yet, the discussion to come is substantially informed by the previous insights.

Notes

- 1 Here I assume that the normative component of thick normative concepts is part of their content. By contrast, pragmatic views hold that the normative information in thick normative concepts is just communicated by some pragmatic mechanism, see, for example, Väyrynen (2013). I will set those views aside, since we are concerned with normative concepts and whether they contain descriptive information.
- 2 For a survey over accounts of thick normative concepts, see Roberts (2011).
- 3 The example originates from Roberts (2011, p. 496).
- 4 For helpful distinctions between versions of perspectivism, see Kieseewetter (2011).
- 5 Jackson (1991, pp. 462–463).
- 6 For a recent account that defends conceptual constraints on moral systems, see Cuneo and Shafer-Landau (2014). The account aims to respond to epistemic challenges raised against the position of non-naturalism.
- 7 This rough characterization suffices for our present purposes. Note, however, that an account of thin normative propositions can be given in terms of normative states and their descriptive parts.
- 8 To name just one example: it has been argued that normative concepts are best construed as thin normative concepts in order to account for genuine normative disagreement, see Hare (1964). For critical discussion, see Plunkett and Sundell (2013).
- 9 See Hare (1964).
- 10 See, for example, McPherson (2018) and references therein. For the importance of distinguishing between authoritative and other normative concepts in metaethical debates, see Plunkett (2020). See also Väyrynen (2021) for criticism of normative contingentism in terms of authoritative normativity. For applications in epistemic debates, see Maguire and Woods (2020).
- 11 See McPherson and Plunkett (2024) for an argument that there are multiple, equally optimally apt candidates for a notion of authoritativeness.
- 12 See Baker (2017, p. 568).
- 13 The application to the practical question of what to do is most straightforward. However, I take the considerations to apply to normativity in general. In the epistemic debate, for example, we are interested in settling what-to-believe-questions.
- 14 McPherson uses the label ‘practical ought’. I will use the more general label.
- 15 McPherson (2018, p. 254). McPherson credits Wedgwood (2004) who in turn credits Cullity and Gaut (1997) for the initial characterization of (Sticky Situation).
- 16 For reasons of convenience, I will drop the parenthesis.
- 17 McPherson (2018, p. 255). In the same way, one might need authoritatively normative concepts to make sense of their denial; e.g., Nietzsche (1998) criticizes the view that morality is normatively significant.
- 18 The application to normative reasons is more controversial. Some have argued that reasons are precisely what separate the authoritatively normative and the formally normative domain. See Parfit (2011, sec. 88). I will address this point in due course.
- 19 See McPherson and Plunkett (2024, p. 5, fn. 16).
- 20 Here I follow the presentation in McPherson (2018).
- 21 See Parfit (2011, sec. 88).
- 22 Scanlon (1998, p. 19).

- 23 [McPherson and Plunkett \(2024, p. 5\)](#).
- 24 The reader might not be convinced that ‘all-things-considered’ must be read norm internal. In this case, both characterizations might simply track the same thing. In Sections ‘Further Suspect Objection’ and ‘Authority and Its Modal Profile’, I argue that it is important that authoritatively normative concepts do not invoke a further standard. I will stay neutral on whether ‘all-things-considered’-characterizations require a standard of consideration. This is why I use the authority label for the rest of the book. For details, see the mentioned sections.
- 25 The proposal that I will defend at the end of this chapter has important similarities with the subset view. However, I disagree with a number of other claims made in Howard and Laskowski (forthcoming).
- 26 See [Hansson \(2014, p. 7\)](#).
- 27 For a reductive account, see [McPherson \(2018\)](#).
- 28 It has to be admitted, though, that it is not obvious how to interpret the relevant locution ‘X is the thing to do’. So, one might ask to what extent it helps to fix the relevant concept. I will assume that it provides us at least with an intuitive grip on the concept. Later, in Section ‘Work for a Theory of Authoritative Normativity’, I will propose a relevance-based approach to elucidate the concept of authoritative normativity. Worries about the coherence of the concept will also be addressed in that section.
- 29 For a precise formulation of how to define trivial application conditions, I have to beg the reader’s patience. The formal explication will be given in Section ‘Authority and Its Modal Profile’ in terms of *tied* parts of states. Even authoritatively normative concepts arguably presuppose the agent’s existence. However, the tied-part analysis will explain the sense in which these constraints are trivial.
- 30 See [Plunkett \(2020\)](#) and [McPherson and Plunkett \(2020\)](#).
- 31 See [McPherson and Plunkett \(2020, p. 288\)](#). Note that concepts can be extension-revealing even if they function like, e.g., operators. Whether such a concept is extension-revealing depends on how much it reveals of the truth-conditions of propositions expressed by sentences that involve the operator. See [McPherson and Plunkett \(2020, p. 288, fn. 50\)](#).
- 32 See [McPherson and Plunkett \(2020, p. 289\)](#). The authors focus on inferences that are ‘relatively straightforward for competent users of these concepts to make’. They leave it open that it happens that concepts that are not authority-revealing are in fact maximally authoritative.
- 33 Note that it does not necessarily help to settle the what-to-do-question correctly. As mentioned in endnote 32, being authority-revealing and being authoritative need not fall together. In the discussion to follow, I will focus on the latter feature of normative concepts.
- 34 [Plunkett \(2020, p. 189\)](#). argues in favor of such a tension. He admits, however, that this might be a contingent fact about us.
- 35 See [Moore \(1903, §13\)](#). For similar considerations, see [Gibbard \(2003\)](#).
- 36 Note that this reasoning parallels an argument in [McPherson \(2012\)](#) made in a metaphysical context. Roughly, McPherson argues that those who assume *sui generis* normative facts face a dilemma. Either normative and natural facts are intimately related so that they turn out to be metaphysically continuous which undermines the special status of normative facts, or they are importantly distinct in which case connections between them are brute. Brute connections,

- however, make it difficult to see why particular combinations obtain, rather than competing combinations.
- 37 [McPherson \(2018\)](#) develops an analysis of AUTHORITATIVE OUGHT in terms of *non-arbitrary* selection. While my argument is based on his insights, it is not committed to his analysis.
- 38 See [McPherson \(2018, §4.1\)](#).
- 39 This strategy is also pursued by McPherson when he requires that AUTHORITATIVE OUGHT is the ought of a norm that is ‘the norm to appeal to’ in the relevant context ([McPherson 2018](#), p. 16).
- 40 Note that pragmatic considerations merely license conclusions about what pragmatics requires us to do.
- 41 For elaborate discussion, see [Clarke-Doane \(2014\)](#).
- 42 The pluralist can still argue that those normative concepts that are socially or psychologically relevant to us are autonomous. But this claim reveals (if anything) more about us than it does about normativity.
- 43 [Baker \(2017, p. 568\)](#).
- 44 Others even doubt that there is problem in search of a solution. Howard and Laskowski (forthcoming) argue that the all-things-considered ought is all we need.
- 45 [Baker \(2018, p. 11\)](#). For ease of presentation, I follow Baker in speaking loosely about different *oughts*.
- 46 One might wonder whether the discussion could be more helpfully phrased in terms of normative reasons. See, for example, the discussion in Howard and Laskowski (forthcoming). However, since my focus here is on discussing how different oughts relate to each other, and I’m skeptical about the notion of an authoritative reason, I will restrict my discussion to oughts-talk. However, the arguments can be translated into reasons-talk by referring to sets of reasons.
- 47 One might object that for this strategy to work, we must posit a privileged relevance ordering, since there are many orderings. Let me point out two things in response. First, the concept of authoritative normativity presupposes that there is a privileged answer to what-to-do-questions. So, we cannot expect our account to get along without this assumption. Second, according to Baker’s objection, authoritative ought is just another ought among many. A normative relevance ordering, however, is put to work once all possible oughts (or accumulation functions) are on the table. They are ranked according to their normative relevance. This constitutes a plausible sense in which we did not just invoke another ought. Note that I’m sympathetic to Howard and Laskowski’s proposal that an all-things-considered ought is the authoritative ought because it is the only ought that takes into account the total set of reasons. In terms of my proposal, this would mean that the all-things-considered ought has always the highest rank in the normative relevance ordering. I discuss this point in more detail in Section ‘Authoritatively Normative States’.
- 48 [Baker \(2018, p. 29\)](#).
- 49 [Baker \(2018, p. 27\)](#).
- 50 See [Schaffer \(2000\)](#).
- 51 For example, [Krämer and Roski \(2017, p. 1207, fn. 29\)](#) outline an objective and a pragmatic way to conceive of their relevance criterion, but stay agnostic about which one should be adopted. Likewise, there are different ways to conceive of a normative relevance ordering. As will see, this is a strength of the account.

- 52 It has been pointed out that the authoritative/non-authoritative distinction marks a categorical difference, whereas relevance is a matter of degree. See [McPherson and Plunkett \(2024, p. 4\)](#). However, it is also well known that a categorical distinction can be defined in terms of gradable concepts. Possessing the highest rank in a relevance ordering defines a categorical distinction.
- 53 One way of specifying the normative relevance ordering is in terms of the weights of sets of reasons. However, the notion of a normative relevance ordering is more general.
- 54 [Baker \(2018, p. 578\)](#).
- 55 [Baker \(2018, p. 578, italics in original\)](#).
- 56 A pre-order is a binary reflexive and transitive relation. I do not assume anti-symmetry because I want to stay neutral concerning the possibility of normative, in particular, authoritatively normative dilemmas. This requires that there can be different authoritatively normative states that possess the same normative relevance. One might doubt that there are such dilemmas given the motivation for authoritative normativity. Yet, since moral dilemmas are considered in the moral debate and morality is a good candidate for a normative system that possesses authority, the account should not rule out this option.
- 57 Recall that I have assumed that non-cognitivist positions can account for the autonomy debate in terms of minimalism with respect to normative propositions and their truth. At this point, I think, the non-cognitivist must tell a slightly different story. The idea that normativity comes with a particular relevance ordering most naturally fits a realist position. However, there is something non-cognitivists can say. Details depend on the accounts, but if they, for example, assume that moral claims express particular attitudes toward the act in question, authoritatively normative claims might additionally express an attitude toward the relevance ordering of those attitudes expressed by moral and prudential claims.
- 58 Another option would be to take the normative systems in *I* to be ordered. Note that this is a special case of the state-ordering. So, in assuming the state-ordering *I* am working with the weaker assumption.
- 59 For simplicity, I restrict the account to fully determined courses of action. The reason is that, for example, obligations that allow for several forms of compliance require a more complicated account of truthmakers of the corresponding propositions. This complication more generally applies to truthmaker-based accounts of normative operators. Since it is orthogonal to the current topic, I will set it aside. For elaborate discussion, see [Faroldi \(2019, §5.3\)](#).
- 60 Note that this is a technical term. It should be taken with some caution because there is also a natural sense in which the agent's existence is decisive for the obtaining of any normative states concerning them.
- 61 A maximally tied part of states is the null-state because it is part of every state.
- 62 One might object that given that authoritatively normative concepts settle what-to-do-questions, AUTHORITATIVE OUGHT must be restricted to possible courses of action. Moreover, since morality is a plausible candidate for a normative system that settles what we authoritatively ought to do and many have assumed the Moral-Ought-implies-Can Principle, we might think that the Authoritative-Ought-implies-Can Principle holds as well. In the state space, we can account for this claim by assuming that $o!^{\phi}$ contains descriptive parts to the effect that ϕ -ing is possible. Yet, one might think that this is incompatible with the only-tied-parts-constraint. In response, we can distinguish two ways for a course of action to be permitted. On the one hand, it might be permitted to ϕ

just because it is impossible not to ϕ . On the other hand, it might be permitted to ϕ though not- ϕ -ing is a live option. Permissions of the latter kind require that ϕ -ing and not- ϕ -ing are possible. Thus, the possibility to ϕ is a tied part.

- 63 Note that depending on our purposes, we can consider conceptual, metaphysical, or some other form of necessity and contingency.
- 64 This is only a preliminary formulation. In [Chapter 8](#), I will invoke a relevance constraint.



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

Part III

**Metaphysical Autonomy
Theses**



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

7 Ground-Based Autonomy¹

Ground Works

A normative autonomy thesis is a claim about how the normative relates to the non-normative. In his *Ethics: Inventing Right and Wrong*, Mackie (1977) raises an important question concerning the nature of that relation:

What is the connection between the natural fact that an action is a piece of deliberate cruelty—say, causing pain just for fun—and the moral fact that it is wrong? It cannot be an entailment, a logical or semantic necessity. Yet it is not merely that the two features occur together. The wrongness must somehow be ‘consequential’ or ‘supervenient’; it is wrong because it is a piece of deliberate cruelty. But just what in the world is signified by this ‘because’?²

The answer to that question that I am going to defend in this chapter is easily given: it’s *grounding*.³ Roughly, grounding is an objective, non-causal priority relation among facts or truths. Some facts or truths seem to be more fundamental than others. Moreover, some facts or truths seem to be derivative because they obtain or hold *in virtue of* more fundamental facts or truths. Consider, for example, an arbitrary conjunctive truth such as [Macron is the French president and male]. This conjunctive truth holds in virtue of its two conjuncts together. That Macron is the French president and male holds partially in virtue of the truth that Macron is the French president and partially in virtue of the truth that Macron is male. Consider another example: suppose that this stone is green. This truth is grounded in the more fundamental truth that this stone is emerald. Grounds make it the case that another truth—the *groundee*—holds. The crucial question in the context of autonomy is whether any normative truth is *fully* grounded in descriptive truths. To pick up on the previous example while assuming that ‘because’ refers to grounding, we want to know whether the action is wrong *fully* because it is a piece of deliberate cruelty.

To begin with, let me elaborate on the notion of *ground*. I hope to stay neutral with respect to most of the ground-internal debates. However, for reasons of convenience, I will make some choices while nothing substantive hinges on them. Moreover, in some debates I have to take a stance, but again, although particular cases might depend on these assumptions, as far as I can see, none of them is crucial for a ground-based autonomy debate. Let me outline the assumptions I am going to make. To introduce the notion of ground, I have spoken of ground as a relation between facts or truths. One debate concerns whether ground is fundamentally a *relation*. It is important to notice that our way of expressing claims of ground is not neutral on that debate. Conveying grounding claims by means of a *predicate* like ‘ground(s)’ that is combined with two or more singular terms of the form ‘the truth that’ or ‘the fact that’ seems to commit us to a relational approach of ground. On the other hand, we can also use *sentential connectives* such as ‘because’ to convey grounding claims which seems to carry no such commitment. It has been put forward that assuming a sentential connective is also ontologically more parsimonious, for it does not presuppose the existence of the relata.⁴ This trait is significant in the metanormative context. Antirealists accept normative sentences, but deny the existence of normative facts. In order to open the autonomy debate for a wide metanormative audience, it might be best to adopt the connective-approach. For ease of exposition, however, I shall often speak of ground as a relation.

Another debate concerns the relata of ground.⁵ We can distinguish between *representational* and *worldly* approaches to ground. While representational approaches take *propositions* to be the relata of the metaphysical priority ordering, worldly approaches consider *facts* instead. This distinction bears on how fine-grained the relata are distinguished. Commonly, propositions are sensitive to distinctions that do not affect the individuation of facts.⁶ For example, the proposition [It is not the case that it is not the case that it is rainy] and the proposition [It is rainy] are plausibly taken to be distinct propositions, while they may represent the same fact. I have already argued that propositions provide a common ground for modal and ground-based autonomy theses that allows for a fruitful comparison. Therefore, I take propositions to be grounded in and to ground other propositions.

I use the terms ‘because’ and ‘in virtue of’ to refer to the grounding relation. I assume that grounding is irreflexive, asymmetric, and transitive.⁷ Accordingly, nothing grounds itself; if one proposition helps to ground another proposition, then the converse does not hold; and grounding can be chained such that it makes sense to talk about mediate grounds. Grounding either is or backs an explanatory connection.⁸ Sometimes, I will simply say that one proposition explains another proposition to state

that the latter is grounded in the former. Grounding is a factive relation. This means that whenever one proposition is grounded in other propositions, grounds and grounded propositions are true. However, there is also a non-factive notion of ground that will become important for the account that I am eventually going to defend. According to the non-factive notion of ground, it is not required that grounds and groundee are true.⁹

Another distinction that is particularly relevant in the context of an autonomy debate concerns the completeness of grounds. We already encountered this distinction in one of the examples. Reconsider the proposition that Macron is the French president and male and its grounds. One partial ground is the proposition that Macron is male. A full ground of the proposition comprises both conjuncts. At this stage, I will stay neutral with respect to any priority between partial and full ground. That is, I will not commit myself to the claim that a partial ground is always part of a full ground.¹⁰ Though, this is true for most cases. This concludes the ground-theoretic preliminaries.

Let me return to the autonomy debate. It has been argued that grounding helps to structure several debates such as the debate about physicalism.¹¹ I propose that the autonomy debate is another one. Yet, like in the logical autonomy debate, a ground-based autonomy thesis requires a careful recognition and treatment of some spurious counterexamples. To illustrate the issue, let me state a preliminary *simple metaphysical autonomy thesis* similar to the simple logical autonomy thesis SLAT:

SMAT: No descriptive propositions fully ground a normative proposition.

Again, SMAT is easily refuted. It is commonly assumed that disjunctions are grounded in their true disjuncts. Thus, descriptive D grounds normative $D \vee N$. So, one might ask about the benefit of a ground-based autonomy thesis. In response, I shall do two things in this chapter. First, I shall argue and explain why an autonomy thesis in terms of ground provides us with a compelling sense in which we cannot ‘get’ something normative from the descriptive. This explains why a ground-based approach improves on a modal approach. Second, I consider an existing contingentist autonomy thesis in terms of ground that aims to deal with spurious counterexamples. I argue that it does not provide a satisfying solution. Together these results motivate a *novel* ground-based explication of autonomy that accommodates for the outlined shortcomings. A comprehensive demonstration of the benefits of that ground-based autonomy thesis has to wait until the next chapter.

Let me conclude this section with a metaphor-based consideration why it is instructive to construe the autonomy thesis in terms of ground. While grounding is often assumed to be primitive which means that it cannot be

defined in more fundamental terms, there is a common telling metaphor: when God made the most fundamental facts, she was done.¹² She did not in addition have to that make the derivative, i.e., grounded, facts. Consider, for example, the fact that this stone is green. Suppose God made the stone emerald. When she was done, she did not make the fact that the stone is green in addition to the fact that the stone is emerald. Now, consider the following reasoning: suppose God made the descriptive facts. Suppose further that there are normative facts. The overarching question of this book is how to understand the claim that we cannot get something normative from the descriptive. If God, in our metaphor, cannot get normative facts from the descriptive facts, she wasn't done when she made the descriptive facts. She had to make normative facts in addition. Thinking in terms of this metaphor shows that a ground-based thesis provides us with a reasonable interpretation of the initial autonomy thesis.

The Why-Not-Whether Argument

In this section, I argue that the claim that the normative is autonomous with respect to the descriptive concerns the question of *why*, not *whether* normative propositions are true. The question of why normative propositions are true in turn concerns the grounds of normative propositions. The contrast between *why*- and *whether*-questions also helps to see why some counterexamples to logical autonomy theses do not pose threats to normative autonomy.

First, and to motivate the contrast, consider a list of ingredients of a particular product. The list tells you *whether* the product contains sugar, but it does not tell you *why* it does. If descriptive premises entail a normative conclusion, the truth of the descriptive propositions necessitates that the normative proposition is true. Thus, if the premises are true, they settle whether the conclusion is true. However, unless the descriptive propositions additionally *ground* the normative proposition, the premises do not explain why the conclusion is true. Take, for example, the conjunction that Macron is the French president and male. It entails and thus settles whether Macron is male, but it does not explain why Macron is male. Now suppose we are concerned with the autonomy of a particular subject-matter-related domain Δ . Suppose further we think that we *can* get Δ from its complement. In that case, we do not only claim that we can answer the question of whether Δ -propositions are true by means of non- Δ -propositions, but also that we can answer the question of *why* Δ -propositions are true by means of non- Δ -propositions.

To develop this claim in more detail, imagine a normatively infallible being. Let's call her HALLOW. Whenever HALLOW makes a moral judgment, the judgment is true. This is not merely a contingent, but an essential

feature of HALLOW. However, HALLOW—and this will be an important contrast case later—is not like God according to divine command theory. HALLOW is infallible with respect to her moral judgments, but she does not *make* moral states of affairs the way they are. Consider the following case: HALLOW judges that Toni ought to keep her promise (in a particular situation). Given HALLOW's essential features, this entails that Toni ought to keep her promise. The fact that Hallow judges that Toni ought to keep her promise settles *whether* Toni ought to keep her promise. Let me scrutinize whether this example undermines the autonomy of the normative domain.

In order to assess that case, one must classify the involved propositions. I assume that the classification of the conclusion as normative is uncontroversial. What is at issue is the classification of the proposition [HALLOW judges that Toni ought to keep her promise]. Recall that I have argued that there is an important difference between normative-belief propositions like [Sarah believes that Toni ought to keep her promise] and normative knowledge-propositions like [Sarah knows that Toni ought to keep her promise]. If I am correct, the former is descriptive while the latter is normative. What explains the difference between the two kinds of propositions is that normative-belief propositions are true iff the subject is in a particular mental state, while normative-knowledge propositions additionally require that the world is in a particular normative state. According to one of the desiderata for a taxonomy, relevantly similar cases should be treated in the same way. So, let me scrutinize whether HALLOW-propositions resemble normative-belief or normative-knowledge propositions. While I shall say a bit more about the classification of HALLOW-propositions at the end of this section, let me make a case for their descriptive status. States to the effect that HALLOW is in a particular mental state suffice to bring about the truth of the proposition [HALLOW judges that Toni ought to keep her promise]. These states only obtain, if the world is in a particular normative state, but this is not part of them. So, *prima facie*, these states do not seem to be normative. Further scrutiny might reveal that the proposition should be classified as normative; for now, let me suppose it is descriptive.

Notice that it would be a surprising result, if the question of whether the normative domain is autonomous would depend on the existence of infallible beings like HALLOW. To the contrary, the two topics seem to be orthogonal. This assessment gets additional support from an assumption that I made throughout this book according to which autonomy is a feature based on the nature of normativity. However, the HALLOW example is a putative counterexample to the autonomy thesis. This tension is solved, if it is recognized that infallible beings (like HALLOW) merely settle *whether* *P* for some suitable *P*. The crucial point, however, is that infallible beings do not settle *why* *P*. That HALLOW judges that Toni ought to keep her promise does not *ground* that Toni ought to keep her promise. Thus,

an autonomy thesis in terms of ground does not depend on the absence of infallible beings.

To lend further support to the conclusion that autonomy concerns why-questions, consider the previously outlined contrast case. According to divine-command theory, the fact that God judges (commands) that Toni ought to keep her promise *makes it the case* that Toni ought to keep her promise. This explanatory connection bears on the autonomy of normativity. Normative truths would be fully generated by non-normative truths. In this scenario, there is a clear sense in which the normative domain is non-autonomous with respect to the domain of God's judgments. That is, God's judgments do not merely settle whether but also why Toni ought to keep her promise. So, God but not HALLOW threatens the autonomy of normativity. This constitutes a way in which a ground-based autonomy thesis is superior to modal construals.

Finally, let me return to the classification of HALLOW-propositions. It seems that, on the assumption that HALLOW-propositions are descriptive, the example refutes the Re-Gap Principle as well as the restricted Re-Gap Principle.¹³ The reason is that normative states that exactly exclude the conclusion would plausibly have to be counted as being wholly relevant to the exclusion of states that make true the HALLOW-proposition. At the same time, this is because in HALLOW's case—in contrast to Sarah's case—moral judgment is *factive*. It is indeed essential to HALLOW that she is morally infallible. This suggests that states that make true propositions of the form [HALLOW judges that *N*] contain some normative part that guarantees that HALLOW tracks morality. Therefore, I think, it is reasonable to conceive of HALLOW-propositions as broadly normative.¹⁴ I admit that considerations along these lines require substantial work. It is not trivial to assess which states make HALLOW-propositions true and there might be different ways to work out the details. To me, however, it seems plausible that HALLOW-propositions differ from ordinary judgment-propositions in being not purely non-normative. Anyway, the general question is whether we got something normative from something descriptive. If one adopts a demanding construal of what it means to 'get' something normative, e.g., in terms of ground, one can reject the example. HALLOW's judgments do not *explain* the normative truth in question. If one instead adopts a more liberal construal, e.g., in modal terms, one might need a more restrictive notion of the descriptive that excludes premises that smuggle in normative presuppositions. This illustrates another strength of a ground-based autonomy thesis.

The Essential Argument

The second argument in favor of a ground-based explication of an autonomy thesis concerns the significance of autonomy. In contrast to Pigden's

(1989) assessment, people have taken the autonomy thesis to be a thesis that is worth a genuine debate. But why does it matter whether the autonomy thesis is true? A ground-based autonomy thesis provides us with the following answer: the truth or falsity of a ground-based autonomy thesis reveals something about normativity, namely what is at stake in assuming normativity. I shall now develop the details of the argument.

To begin with, consider a contrast employed in Schaffer (2009). Schaffer contrasts what he calls the *Quinean View* and the *Aristotelian View*. According to the Quinean View, the basic metaphysical questions are *existence* questions like the question of whether numbers or properties exist. By contrast, the Aristotelian View is concerned with the question of *how* things exist, whether their existence is grounded, and what grounds it.¹⁵ It focusses on the question of whether entities like numbers are derivative or fundamental. Schaffer points out that on the Aristotelian View, there is ‘no longer any harm in positing an abundant roster of existents, *provided it is grounded on a sparse basis*’.¹⁶

Let me explain how this contrast bears on the autonomy debate.¹⁷ Recall SMAT: no descriptive propositions fully ground a normative proposition. Suppose first that the contrary is true: *every* normative proposition is fully grounded in descriptive propositions. If this were true, the Aristotelian question has an answer, namely that the normative is derivative. If Schaffer is correct, then this result diminishes any harm in positing normative truths. Now, suppose that some ground-based autonomy thesis is correct. As a result, the normative is not completely derivative. Thus, positing normativity is not that harmless. Normativity does not fully depend on a sparse—in our case non-normative—basis.

The previous considerations reveal that a ground-based autonomy thesis is a claim about the nature of normativity. The connection between a ground-based autonomy thesis and the nature of normativity is robust because grounding requires an intimate connection between grounds and groundee. In order to illustrate, consider meta-grounding views according to which grounding truths of the form X is grounded in Γ are grounded in the essences of the involved items.¹⁸ For example, that [Socrates exists] grounds [Singleton Socrates exists] is arguably no coincidence. According to essence-based approaches, it lies in the nature of the singleton comprising Socrates and only Socrates to be intimately connected to Socrates’ existence. Note, however, that no commitment to essence-based approaches is required. On alternative approaches that assume, for example, some *basic* grounding principles that link grounds to the groundee, these basic principles plausibly have an equal claim to settle fundamental traits of the normative.¹⁹ The close connection between questions concerning autonomy and the nature of the normative reveal why a ground-based autonomy debate is worth caring about and why it requires a genuine debate.

One consequence of this explication of the autonomy thesis is that it is not necessarily unique to the normative domain, but based on its characteristic features. There might be other domains that are autonomous in the ground-theoretic sense. If, for example, universal propositions are not fully grounded in particular propositions, the universal domain would turn out to be autonomous in the explicated sense. Thus, autonomy need not be *unique* to the normative domain. However, each ground-based autonomy thesis, e.g., a normative or a universal autonomy thesis, tracks a feature that is inherent in the domain in question. It is a characteristic feature of the *normative* domain what grounds it and whether normative propositions are fully derivative. Similarly, it is a characteristic feature of the universal domain what grounds it and whether universal propositions are fully derivative. I take the close connection to crucial features of the domain in question to be another reason to favor a ground-based autonomy thesis.

This concludes my preliminary motivation of a ground-theoretic turn in the autonomy debate. It will be instructive to see how a ground-based autonomy debate relates to one of the debates to which the notion of ground has been applied in the interest of structuring it, namely the debate between naturalists and non-naturalists.²⁰

Non-Naturalism and Normative Autonomy

Roughly, non-naturalism is the combination of a realist view according to which there are normative facts and the view that these facts are *sui generis*, i.e., of their own kind. Naturalists agree with the first part, but deny the second. Rosen (2017) suggests the following ground-based explications of naturalism and non-naturalism, respectively.

Ethical Naturalism:	For all normative properties F , for all x : if Fx then there exists a non-normative condition φ such that Fx is metaphysically grounded in $\varphi(x)$. ²¹
Bridge Law Non-Naturalism:	Every normative fact of the form $F\alpha$ is grounded in some non-normative fact $\varphi(\alpha)$, together with a bridge law connecting F and φ . ²²

It is easily verified that if Ethical Naturalism is true, SMAT is false. By contrast, Bridge Law Non-Naturalism dovetails with SMAT given the plausible assumption that bridge laws are normative. Thus, one might suggest that a ground-based autonomy debate reduces to the debate between naturalists and non-naturalists.²³ Indeed, some have claimed that all non-naturalist views ‘share the core idea that morality is autonomous’.²⁴ In the remainder of this section, I argue that the ground-based autonomy debate

does not reduce the naturalism–non-naturalism distinction. For one thing, a ground-based autonomy debate matters for metanormative views beyond that distinction; for another thing, a ground-based autonomy debate improves distinctions between different forms of non-naturalism.²⁵

Let me first explain how the ground-based autonomy debate can shed light on the assessment of alternative realist views, specifically non-naturalist views. Consider the following realist view: there are normative facts (the realist part); the normative facts are (i) *sui generis* and (ii) fully grounded in descriptive facts. This view apparently combines features of naturalism and non-naturalism. According to this view, normative facts are *sui generis* which gives the account a non-naturalist flavor. On the other hand, the account assumes that normative facts are fully grounded in descriptive facts. Consequently, they are fully derivative. This consequence in turn seems to push the account toward the naturalist camp.²⁶ Indeed, according to the ground-based explication of naturalism, the account is classified as a naturalist view. The outlined realist view assumes what Rosen (2017) calls *Moorean connections*. We face a Moorean connection iff for some normative property *N*: (i) something is *N* just in virtue of some non-normative *N*-making feature and (ii) *N* is *sui generis*.²⁷ One might insist that Moorean connections are implausible. This might very well be true, but that does not settle whether it is an (implausible) naturalist or non-naturalist view.

I am now in a position to explain how the autonomy debate clarifies the assessment of that view. The autonomy debate helps to solve the tension between arguments according to which characterizations of non-naturalism must be able to account for Moorean views and arguments according to which Moorean views already grant too much toward naturalists.²⁸ We can grant that Moorean views fall on the non-naturalist side of the divide because they take normative facts to be *sui generis*. They share this assumption with, for example, Bridge-Law Non-Naturalists and it distinguishes them from naturalists. However, Moorean views fall on the *non-autonomist* side of the divide. This explains why they share features with naturalism. Moorean and naturalist views assume that normativity is not autonomous. This distinguishes them from other metanormative views, particularly from other non-naturalist views, that assume that normativity is autonomous. Consequently, the Moorean view is a non-autonomist non-naturalist position which explains why it seems to be an intermediate position between naturalism and non-naturalism.

Next, I show that the ground-based autonomy debate concerns views beyond the realist camp. The autonomy debate thus exceeds the naturalism–non-naturalism debate. I have already argued that views that accept a minimal notion of normative truths can take part in the modal autonomy debate. In particular, all views that aim to account for complex normative

judgments such as if-then-statements and notably for normative reasoning must provide some account of normative truths. To give just one brief example, consider a view according to which normative judgments such as ‘Torture is morally wrong’ express disapproval toward torture. Such an account should be ready to give an explanation for the meaning of sentences embedding the former sentence like ‘I wonder whether torture is morally wrong’ or ‘If torture is morally wrong so is getting someone else to torture someone’. These sentences cannot be explained just by pointing to disapproval toward torture. This problem is called the *embedding problem* or the *Frege-Geach problem*.²⁹

It is beyond the scope of this book to discuss the embedding problem or approaches to respond to it in detail.³⁰ For my current purposes it is important, however, that metanormative views that accept the challenge raised by the embedding problem such as hybrid forms of expressivism or quasi realism are concerned with the autonomy debate. Schroeder (2009) discusses a *modal* case that illustrates the significance of autonomy questions for those views: from ‘Jeremy said that stealing is wrong’ and ‘Everything that Jeremy said is true’ we can infer ‘Stealing is wrong’. Schroeder argues that to secure the argument’s validity, hybrid forms of expressivism are committed to the claim that ‘Everything that Jeremy said is true’ expresses a desire-like attitude.³¹ On the hybrid view, coming to accept a normative conclusion is to accept both a belief and a desire-like attitude. The premises must not only guarantee the descriptive belief-component but also the normative desire-like component. Schroeder assumes that ‘Jeremy said that stealing is wrong’ is uncontroversially descriptive. He concludes that on the hybrid view, ‘Everything that Jeremy said is true’ turns out to be normative. If Schroeder is correct about this, on the hybrid view, the example is compatible with modal autonomy.

When we are concerned with a *ground*-based autonomy thesis, one might think that its scope should be restricted to realist positions. For concreteness, consider the claim ‘Sarah ought to help her grandmother *because* she promised her to help’. One reading of that claim presupposes a realist position: the normative fact that Sarah ought to help her grandmother obtains in virtue of the descriptive fact that Sarah promised her to help. The connective approach to ground, however, does not presuppose a realist position. Let me illustrate that the assessment of because-claims such as ‘*N because D*’ matters to, e.g., quasi realism. In a recent paper, Berker (2020) discusses how quasi realists can do justice to one of their central aims, namely the mind-independence of normative matters, without reneging the commitments of their accounts according to which normative sentences express a particular state of mind, e.g., disapproval. According to Berker, quasi realists need to be able to sincerely utter sentences like ‘It is not the case that (kicking dogs for fun is wrong because we disapprove of it)’. As soon as these views provide an account of the

relevant because-claims the question that is scrutinized in this book arises. Can because-claims of the form ‘*N because D*’ be true if they are understood as *complete* explanations? I think that answering this question is important for every account of normativity. Moreover, as far as I can see, antirealist accounts are not committed to a particular answer.

Let me take stock: the autonomy debate is partly independent from the naturalism–non-naturalism-debate. Moreover, taking position in the metanormative terrain does not in general imply a particular position in the autonomy debate and vice versa. This illustrates the significance of a satisfying ground-based approach of autonomy. It provides another dimension along which the metanormative terrain can be structured. In the next section, I shall briefly illustrate why a simple ground-based thesis like SMAT is not satisfying. Against this background, I discuss a ground-based autonomy thesis from the literature that provides a contingentist response to the outlined issues.

A Preliminary and a Contingentist Approach

Not all examples that threaten logical autonomy theses matter to a ground-based thesis. Considerations in the context of the why-not-whether argument revealed that examples like the HALLOW-example do not get a foothold on ground-based autonomy theses. The same is true for (Truth-Teller). However, it is important to bear in mind that disjunctive cases refute the simple ground-based autonomy thesis SMAT. Mixed disjunctions with a true descriptive disjunct provide counterexamples to SMAT. [Tea-drinking is common in England] grounds [Tea-drinking is common in England or all New Zealanders ought to be shot]. Moreover, several spurious counterexamples from the logical autonomy debate transfer to the ground-based debate: [Mary does not exist] grounds [There is no obligation for Mary to ϕ]. In response, Maguire (2015) proposes a ground-based autonomy thesis that aims to accommodate for such cases. Maguire takes the focus on logical autonomy theses to be a mistake. He argues that a ground-based autonomy thesis can explain the success of counterexamples to logical autonomy theses.³² First, I present his account that is based on facts having a contingent taxonomic status. I argue that Maguire’s autonomy thesis is subject to serious objections that render it less convincing as a substantial autonomy thesis. I introduce worries both concerning the taxonomic consequences and the informative value of Maguire’s thesis.

Two Roles of Ground in the Autonomy Debate

Maguire’s autonomy thesis focuses on facts which he takes to be worldly items.³³ Note, however, that on Maguire’s account, facts are taken to be true propositions. One important thing to notice about the way Maguire

presents his autonomy thesis is that he does not address the two questions that concern the taxonomic status of facts and the way *autonomy* is understood separately. He explicitly states that he is concerned with a *closure principle*, i.e., with the question of which operations and relations *preserve* the taxonomic status of facts.³⁴ While this strategy is not worrisome in itself, it will become clear soon that it invites the worries that I will point out. Maguire's autonomy thesis includes two separate, complementary parts:³⁵

Metaphysical Autonomy (MA):	No normative fact is grounded just by descriptive facts.
Converse Metaphysical Autonomy (CMA):	Any fact partly ³⁶ grounded by a normative fact is a normative fact.

Maguire's autonomy thesis does not presuppose an independent taxonomy of facts. However, he assumes an intuitive grasp on paradigmatic cases of both descriptive and normative facts as I did throughout the book. His account is supposed to determine the taxonomic status of non-paradigmatic cases. The key assumption is that grounding preserves taxonomic status. I will discuss the disjunctive case to illustrate how Maguire's autonomy thesis works. In the rest of this chapter, I will use the square-brackets notation to refer to facts instead of propositions. Recall that Maguire's account applies just to facts, i.e., true propositions. Assume $[D]$ and $[N]$ are paradigmatic cases of a descriptive and a normative fact, respectively. According to common assumptions about grounding disjunctions, there are three schematic ways to ground the disjunctive fact $[D \vee N]$: (i) $[N]$ is the sole ground, (ii) $[D]$ is the sole ground, and (iii) $[D]$ and $[N]$ together ground the disjunctive fact.³⁷ In the first case, since $[N]$ is normative $[D \vee N]$ is normative by CMA, for grounding preserves taxonomic status. The example is compatible with MA because the sole ground of $[D \vee N]$ is normative. Thus, a normative fact is grounded in a normative fact. The third case is similar since here $[D \vee N]$ is *partly* grounded in $[N]$. So again, $[D \vee N]$ is normative by CMA and compatible with MA because one of its grounds is normative. In the second case, (ii), $[D \vee N]$ is descriptive because it is grounded just by the descriptive fact $[D]$ and hence, according to MA, cannot be normative. Thus, a descriptive fact grounds a descriptive fact. This result is compatible with MA. By distinguishing these cases, Maguire can show that the putative counterexample does not refute his autonomy thesis. This result is due to the assumption that the taxonomic status of facts is contingent.

Maguire's explication of the autonomy thesis employs grounding to settle two questions at once: on the one hand, grounding is used to explicate

what autonomy consists in. The normative domain is autonomous with respect to the descriptive domain in virtue of the fact that normative facts are not *grounded* just in descriptive facts. On the other hand, grounding is used to determine the taxonomic status of non-paradigmatic cases like mixed disjunctive facts. In the first chapter of this book, I have argued in favor of a methodological constraint on taxonomies that suggests that our taxonomic principles should not presuppose autonomy principles. In what follows, I shall argue that the way Maguire employs the notion of ground violates this constraint.

Too Much Work for a Notion of Ground

Maguire's ground-based autonomy thesis incorporates the function of specifying the autonomy thesis and the descriptive-normative distinction. First, I raise a problem for the resulting classification of facts. Thereafter, I raise some doubts concerning the substantiality of Maguire's autonomy thesis.

Let me start with what I shall call the *Problem of Contingentism*. Most notably, Maguire's autonomy thesis entails that the taxonomic status of facts (more precisely, instances of fact-types) is contingent and depends on the actual grounds of the fact. We already saw that the disjunctive fact $[D \vee N]$ can be normative or descriptive depending on whether $[N]$ obtains. It also can be both normative and descriptive if both disjuncts obtain.³⁸ Consequently, Maguire's autonomy thesis is committed to *contingentism* concerning the taxonomic status of facts.³⁹ Now, the worry is that contingentism entails implausible taxonomic consequences. Recall that Maguire assumes that at least in some cases we have an intuitive grasp of the taxonomic status of facts. Consider the following two disjunctive facts: [It is raining or murder is morally wrong] and [It is raining or charity is morally wrong]. If it rains and murder is wrong while charity is not, then the first fact is normative while the second one is not. Granted that the taxonomic status of disjunctive facts is not trivially accessible, it still seems to be implausible that the two facts belong to different categories concerning their normative status. Both disjunctive facts seem to be partly about rainy weather and partly about moral wrongness. Since Maguire's account relies on a difference in subject matter between the paradigmatic cases, denying that the status of disjunctive facts depends on what they are about is questionable.

Note that this taxonomic result seems to violate the *Consistency-desideratum* of a taxonomy, according to which relevantly similar cases are treated in the same way. However, I emphasized that one way to satisfy the desideratum is to provide a convincing story about why certain cases might just appear to be similar, while in fact being relevantly different. So,

let me scrutinize whether there is such a convincing story to be told. Is there a reason to treat [It is raining or murder is morally wrong] and [It is raining or charity is morally wrong] differently in a context where it is rainy, murder is wrong, and charity is not wrong? The only reason I can see is that the former but not the latter obtains in virtue of, depends on, is determined by (add your favorite variation) a normative fact. This, however, is part of what is at issue: are normative facts fully grounded in descriptive facts? I return to this worry in a moment when I put MA in its role as an autonomy thesis under scrutiny. For now, let me focus on taxonomic consequences. Maguire admits that we sometimes cannot know the taxonomic status of a fact until we know its grounds.⁴⁰ This result seems to be particularly unsatisfying given many different questions discussed in metanormative inquiry such as whether normative facts supervene on descriptive facts that for a long time weren't even aware of grounding.⁴¹

Another aspect concerns non-obtaining facts, i.e., false propositions. They are not subject to Maguire's autonomy thesis at all. It is an implausible consequence of the view that it does not classify false propositions. For one thing, false propositions seem to be about a subject matter.⁴² For another thing, it is an implausible consequence of the view that *all* normative propositions are true. Maguire discusses this issue and objects that his account is supposed to be concerned with facts rather than propositions in general. Yet, we took the autonomy thesis to concern the nature of normativity and hence the nature of, for example, normative concepts or properties. The result that the taxonomy is dependent on the actual truth of propositions does not seem to be well motivated and moreover does not seem to fit the initial motivation of the autonomy thesis.

One option to avoid these unintuitive results and to offer a more general account is to understand MA and CMA in terms of non-factive grounding. This is to say that the grounds and the groundee need not obtain in order to correspond to true grounding claims.⁴³ If MA and CMA are understood along these lines, [It is raining or murder is morally wrong] and [It is raining or charity is morally wrong] would both be normative according to CMA. However, note that we would have a normative proposition that is fully grounded in a descriptive proposition. Maguire himself mentions what he calls counterfactual grounding and offers different versions of a general taxonomy of propositions in terms of counterfactual grounding claims. Since he does not adopt one of them as his basic account, I will not discuss these modifications in detail here.

The second worry concerns the substantiality of the resulting autonomy thesis given the double role of grounding. I shall call this objection *Problem of Triviality*. I argue that MA and CMA are together insensitive to what seem to be genuine counterexamples to a ground-based autonomy thesis and moreover are compatible with certain forms of naturalism. To

see why, recall that the disjunctive example, where $[D \vee N]$ is grounded solely by descriptive $[D]$, is a spurious counterexample. Maguire's autonomy thesis is not falsified by the example because $[D \vee N]$ counts as descriptive in those cases. The question arises why one should accept this result. One answer would be that it must be descriptive in order to make this example compatible with the autonomy thesis. This answer would be *ad hoc*. Moreover, it would render the metaphysical autonomy thesis trivially true.

The answer that Maguire seems to suggest and that I previously anticipated is that grounding preserves the taxonomic status.⁴⁴ However, we will see that the preservation principle is incompatible with some forms of non-reductive naturalism. For non-reductive naturalists might want to say that normative facts are fully grounded in (and only in) descriptive facts, and yet normative. In order for the autonomy thesis to be substantial the preserving principle needs to be motivated independently. This can perhaps be done by referring to a common way to articulate grounding claims: the groundee is *nothing over and above* its grounds. In this case, $[D \vee N]$ would be nothing over and above $[D]$. This explanation is independent from the autonomy thesis. The plausibility of the explanation depends on its generality. I will argue that if the explanation generalizes, then Maguire's autonomy thesis is insensitive to genuine counterexamples and, more generally, compatible with a naturalist view of, for example, act utilitarianism. Thus, Maguire's autonomy thesis entails that particular naturalists are autonomists. This, however, seems to be incompatible with a naturalist view. Assume that the fact that act a maximizes overall pleasure grounds the fact that a is morally obligatory. I assume that there are no other grounds for the fact that a is morally obligatory. According to the preservation principle, this fact counts as descriptive. Therefore, the example is compatible with Maguire's autonomy thesis.

Maguire would probably object that this fact is a paradigmatic normative fact and that his autonomy thesis should therefore be rejected in that case. However, I cannot see why paradigmatic examples of normative facts should be excluded from the preservation principle, if grounding preserves the taxonomic status in a way that is independent of Maguire's autonomy thesis. This might not be a decisive reason to reject Maguire's thesis as a form of autonomy thesis. The compatibility results from the trivial truth of the autonomy thesis in those cases. If the naturalist view and the preservation principle are correct, then there are no particular normative facts. Hence, Maguire's autonomy thesis is trivially true. This result, however, fails to do justice to views that deny the autonomy of normativity.

Before taking stock, let me mention one further unsatisfying feature of MA. It states that normative facts are not *just* grounded in descriptive facts. Yet, this does not exclude that they are *fully* grounded in descriptive facts.

I already discussed an example: if $[D \vee N]$ is grounded by $[D]$ and $[N]$, then $[D \vee N]$ is normative. However, it is also *fully grounded* by a descriptive fact, namely $[D]$. So, MA specifies a fairly weak sense of autonomy.

Let me sum up the results. Maguire's autonomy thesis entails implausible taxonomic consequences. It treats relevantly similar cases different. MA and CMA handle spurious counterexamples either in an *ad hoc* way that trivializes the autonomy thesis or by appeal to a more general preservation principle that renders the autonomy thesis insensitive to genuine counterexamples and views that arguably deny ground-based autonomy.

Summary

My aim in this chapter was to motivate a ground-theoretic turn. I have argued that a ground-based autonomy thesis reveals something about essential features of normativity. In contrast to modal autonomy theses, it does not address *whether* normative propositions are true given descriptive premises, but *why* normative propositions are true. However, spurious counterexamples from the logical autonomy debate show that a simple ground-based autonomy thesis is false. Disjunctive counterexamples and in general examples where the grounded proposition is weakly normative threaten a ground-based autonomy thesis. In response, I have discussed a modified variant of the simple ground-based thesis. The proposal assumes taxonomic contingentism. I have argued that the double function of grounding to determine the taxonomic status of non-paradigmatic facts and to explicate a sense of autonomy is objectionable. The resulting taxonomy entails implausible classifications. Moreover, an autonomy thesis that is based on the assumption that grounding preserves taxonomic status is no longer sensitive to genuine counterexamples and metanormative views that deny autonomy. After emphasizing both the advantages of a ground-based thesis and problems for existing approaches, I shall now turn to my own proposal. In a first step, I will condense insights from the logical debate and from the discussion of the contingentist approach in order to formulate desiderata for a ground-based autonomy thesis. Thereafter, I provide a ground-based autonomy thesis that satisfies the desiderata.

Notes

- 1 Chapters 7 and 8 rely on an article previously published by *Synthese*: 'No Normative Free Lunch: Relevance and the Autonomy of the Normative Domain' (2021), vol. 199: 13163–13186. I gratefully acknowledge their permission to reuse this material.
- 2 Mackie (1977, p. 41).
- 3 For influential arguments in favor of the significance of a notion of ground in the (meta-)normative debate, see Berker (2018).

- 4 See [Fine \(2012c\)](#).
- 5 Note that similar issues about the individuation of grounds and groundee matter to connective approaches.
- 6 See, e.g., [Correia \(2010\)](#).
- 7 These are common assumptions, but none of them is uncontested. See, e.g., [Fine \(2010\)](#) and [Schaffer \(2012\)](#).
- 8 See, e.g., [Schaffer \(2016\)](#). For critical discussion, see, e.g., [Maurin \(2019\)](#).
- 9 For discussion of the issue and critical comments on the question of whether one can be defined in terms of the other, see [Fine \(2012b\)](#).
- 10 See [Fine \(2012b\)](#). For elaborate discussion and an account that rejects the priority of full ground, see, e.g., [Trogon and Witmer \(2021\)](#).
- 11 See, e.g., [Dasgupta \(2014\)](#).
- 12 There is an ongoing debate about how to define the fundamental level and whether there are fundamental facts or truths. I cannot go into the details of the debate here, but one natural suggestion would be to think of fundamental truths as ungrounded truths. Later, we will return to the question of fundamentality. For helpful distinctions, see [Leuenberger \(2020\)](#).
- 13 Note that the focus on authoritative normativity is orthogonal to the issue raised by these examples.
- 14 One might wish to reject that HALLOW-propositions are in this way normative. However, this makes room for a different kind of response. If states that make HALLOW-propositions true are purely descriptive, it seems less plausible that the normative states under consideration *exactly* exclude, i.e., are wholly relevant to, those descriptive states. By light of this reasoning, the example is captured by the exception-clause of the Re-Gap Principle.
- 15 See [Schaffer \(2009\)](#), pp. 348–354).
- 16 [Schaffer \(2009\)](#), p. 353, italics in original). Let me flag that I do not agree with Schaffer's way of thinking about grounding as relating objects. I side with views that take existential dependence to be explained in terms of grounding relations among corresponding facts or truths, see, e.g., [Schnieder \(2020\)](#).
- 17 In their recent book *The Moral Universe*, [Bengson et al. \(2024\)](#) argue that the notion of ground is an important ingredient in making progress in recent metaethical debates. I will later discuss their construal of non-naturalism in more detail.
- 18 Essence accounts are discussed in [Rosen \(2010\)](#), [Fine \(2012b\)](#), and [Dasgupta \(2014\)](#).
- 19 See, e.g., [Rosen \(2010\)](#), p. 133).
- 20 See [Rosen \(2017\)](#) and [Bengson et al. \(2024\)](#).
- 21 [Rosen \(2017\)](#), p. 157).
- 22 [Rosen \(2017\)](#), p. 163).
- 23 Pigden discusses what he calls an *ontological* autonomy thesis according to which there are irreducibly moral facts or properties ([Pigden 1989](#), p. 128). He claims that naturalism can be understood as the converse of ontological autonomy.
- 24 [Bengson et al. \(2024\)](#), p. 5).
- 25 So, I disagree with [Bengson et al. \(2024\)](#) that all non-naturalists are autonomists. It is important to note, however, that we are relying on different notions of (moral) autonomy. On their account, the autonomy of a domain consists in the fact that a thorough accounting for its complement will inevitably leave something out ([Bengson et al. 2024](#), p. 65).
- 26 [Leary \(2021\)](#) argues against that assessment. Similarly, [Bengson et al. \(2024\)](#), p. 73) characterize non-naturalism in way that is compatible with normative

- facts being fully grounded in non-normative facts. However, see [FitzPatrick \(2008\)](#) for the view that Moorean views already grant too much toward naturalists.
- 27 See [Rosen \(2017\)](#), pp. 132–133). On Rosen’s account that *N* is *sui generis* means that *N*’s nature does not involve non-normative properties. However, I take the class of Moorean views to be broader, which is compatible with the formulation given in the main text. [Bengson et al. \(2024\)](#), p. 181) accept the conditions (i) and (ii)—and thus count as Moorean view as I understand it—but on their account *N*’s nature might involve non-normative properties.
- 28 Both arguments have been defended in the literature, see Endnote 26.
- 29 Geach called the following thought *the Frege Point*: ‘A thought may have just the same content whether you assent to its truth or not; a proposition may occur in discourse now asserted, now unasserted, and yet be recognizably the same proposition’ [Geach \(1965\)](#), p. 449).
- 30 For an overview, see [van Roojen \(2018\)](#).
- 31 See [Schroeder \(2009\)](#), p. 268, 291).
- 32 See [Maguire \(2015\)](#), pp. 188–189).
- 33 See [Maguire \(2015\)](#), p. 196).
- 34 See [Maguire \(2015\)](#), p. 190).
- 35 Maguire (2015, pp. 194–195). Maguire uses a different terminology. He refers to ethical and non-ethical facts instead. He distinguishes between non-ethical facts and facts that are not ethical. According to his terminology, the domain of non-ethical facts and ethical facts might overlap ([Maguire 2015](#), p. 200, fn. 26). To avoid a misleading terminology, I will refer to descriptive facts instead of non-ethical facts and to normative facts instead of ethical facts.
- 36 Maguire assumes the orthodox definition of partial ground: A fact α partly grounds a fact β iff there is a collection of facts Γ including α that grounds β .
- 37 This is standardly assumed and entailed by [Fine’s \(2012a\)](#) elimination rules for ground.
- 38 At least, this is what Maguire seems to be sympathetic to [Maguire \(2015\)](#), p. 200, fn. 26).
- 39 Recall that I have argued against contingentism in [Chapter 2](#). In what follows, I will make use of some of the considerations that I outlined there.
- 40 [Maguire \(2015\)](#), p. 206).
- 41 Even though grounding might be part of metanormative and normative tradition for longer than usually suggested, see [Berker \(2018\)](#).
- 42 A similar objection is raised by [Maitzen \(2010\)](#): since we do not know whether Goldbach’s Conjecture is true, we do not know the taxonomic status of those conjunctions and disjunctions that involve it.
- 43 See [Fine \(2012a\)](#) for problems of defining non-factive grounding in terms of factive grounding.
- 44 See [Maguire \(2015\)](#), p. 190).

8 Relevant Autonomy Thesis

The Desiderata

In this section, I specify five desiderata for a ground-based autonomy thesis. The desiderata are largely based on insights gained in my previous discussion of logical autonomy theses in [Chapters 4 and 5](#). In this respect, they are not novel. However, it is worth explicitly stating them again in the context of a ground-theoretic discussion because they are part of what motivates a focus on ground. The discussion in [Chapter 4](#) revealed that the autonomy thesis as explicated by existing modal accounts easily becomes an overly weak claim. Any two syntactically or semantically separable domains turn out to be autonomous on these accounts. Moreover, some modal explications are in danger of being acceptable for both autonomists and non-autonomists alike. Consequently, these autonomy theses entail rather insignificant distinctions. The aim to formulate a substantial and distinctively *normative* autonomy thesis motivated two desiderata:¹

- D1:** An autonomy thesis explicates an interesting sense of autonomy that reveals something distinctive about normativity.
- D2:** An autonomy thesis draws a line between spurious and genuine counterexamples and explains why the former do not pose a threat to normative autonomy.

Let me start with the first desideratum. I have argued that a ground-based approach is particularly promising because claims about what grounds normative facts also concern aspects of the nature of normativity. This allows me to specify **D1** as follows:

- G-D1:** An autonomy thesis explicates a sense of autonomy such that autonomy is a feature of the *nature of normativity*.

The second desideratum suggests that some putative counterexamples—the genuine ones—pose a threat to normative autonomy. This reflects that there is a genuine debate about the autonomy of the normative domain. Therefore, it seems reasonable to require that the autonomy thesis entails a non-trivial divide of the metanormative terrain. First, recall that naturalist views are arguably incompatible with an autonomist position. If some descriptive facts, for example, natural facts suffice to fully ground the normative domain, the normative domain would turn out to be a derivative of the descriptive domain and thereby non-autonomous.² This holds in particular, if naturalist views are augmented with a reductivist tendency. Consequently, the autonomy thesis must be false if these views turn out to be correct; they endorse genuine counterexamples. However, an autonomy thesis is not supposed to be an explication of an individual metanormative view. For concreteness, recall the Bridge Law Non-Naturalist as specified in [Rosen \(2017\)](#). The Bridge Law Non-Naturalist is plausibly a proponent of an autonomy thesis. However, an autonomy thesis should not commit autonomists to this particular metanormative view. These observations suggest two additional desiderata for a ground-based autonomy thesis:

- G-D2:** An autonomy thesis is false if certain metanormative views according to which the normative domain is identical with or reducible to the descriptive domain are true.
- G-D3:** An autonomy thesis entails a non-trivial divide of the metanormative terrain.

These desiderata already specify a particular interest. An autonomy thesis that satisfies the desiderata differs significantly from explications of the autonomy thesis that capture a form of merely formal autonomy.³ The autonomy thesis is falsifiable and thus captures more accurately that there is a debate about normative autonomy.

Additionally, **D2** comprises two desiderata that are operative throughout the whole book. An autonomy thesis must be immune against spurious counterexamples and sensitive to genuine counterexamples. So, additionally—though not novelly—I require:

- G-D4:** An autonomy thesis is *not* falsified by spurious counterexamples.
- G-D5:** An autonomy thesis *is* sensitive to, i.e., falsifiable by, genuine counterexamples.

Recall that some spurious counterexamples that originate from the logical autonomy debate matter to a simple ground-based autonomy thesis. Mixed disjunctions like $D \vee N$ are grounded in their true descriptive disjunct D . It is important to acknowledge that our account of spurious

counterexamples according to which the relation at issue does not hold *in virtue of* the normative content is applicable. The ground-based autonomy thesis that I am going to propose handles spurious counterexamples by taking into account only relevant connections between the descriptive and the normative domain. So, the agenda for this chapter is to develop a ground-based autonomy thesis that satisfies the outlined desiderata, G-D1–5.

The Relevant Autonomy Thesis

Spurious counterexamples like the disjunctive case show that an autonomy thesis must be formulated in terms of normative relevance. The mechanism behind several spurious counterexamples exploits features of disjunctions. In general, it exploits features of weakly normative propositions. There is at least one purely descriptive way for the normative proposition to be true. These examples generate issues not only for modal but also for ground-based autonomy theses. The descriptive propositions that do the grounding work are *relevant only* to the descriptive ways for the normative proposition to be true; they are *irrelevant* to those ways for the normative proposition to be true that are responsible for its *normative status*. This motivates the following *relevant autonomy thesis*:

RAT: No normative proposition is fully grounded⁴ in descriptive propositions in a normatively relevant way.

The relevant autonomy thesis is formulated in terms of *normatively relevant grounding*. The next step is to make this concept formally precise. Like in the modal case, the approach is based on truthmaker semantics. In contrast to modal entailment, grounding ensures that the grounds are relevant to the grounded proposition *in some way*, while it does not ensure that the grounds are relevant to the grounded proposition in a *normative way*, i.e., to the *normative parts* of the grounded proposition. Accordingly, one must identify those parts⁵ of a proposition that are responsible for its normative status. One may take the part of a normative proposition N that is responsible for N 's normative status to be all its normative truthmakers. Recall that S^N is the set of all normative states. So, we consider the intersection of the normative proposition N and the set of all normative states S^N : $N|_{S^N} := N \cap S^N$. Intuitively, these are all the normative ways for N to be true. I will use the term *normative realization* to refer to $N|_{S^N}$. I now have the means to make the concept of normatively relevant grounding precise:

Normatively P_1, P_2, \dots fully ground Q in a normatively relevant way
Relevant Ground: iff some of P_1, P_2, \dots fully ground $Q|_{S^N}$.

Note that normatively relevant grounding is not supposed to be a different kind of grounding. The relevance constraint simply excludes the possibility of certain sorts of grounding claims.⁶ With this concept of normatively relevant grounding at hand, I can make the RAT on the propositional level formally precise:

Prop RAT: No normative proposition Q is such that its normative realization $Q|_{SN}$ is fully grounded in descriptive propositions.

Since the present account is based on truthmaker semantics, it is worth explicitly stating the thesis that corresponds to the RAT on the state level.⁷ To this end, I need truthmaker semantics for statements of ground. A promising candidate is [Fine's \(2012c\)](#) approach. He defines the notion of weak full ground as follows:

Weak Full Ground (\leq): P_1, P_2, \dots weakly fully ground Q iff $s_1 \sqcup s_2 \sqcup \dots$ makes Q true whenever s_1, s_2, \dots make P_1, P_2, \dots true, respectively.⁸

Note that, in contrast to our more intuitive understanding of ground, weak full ground is *reflexive*. This is, for any P we have $P \leq P$.⁹ One might object that the RAT thus explicated is trivially true. To see why, note that if I adopt Fine's truthmaker semantics for ground and assume that no normative state is identical to a descriptive state, then by definition it is impossible for a descriptive proposition, i.e., a set of descriptive states, to ground a normative realization of a proposition, i.e., the maximal normative subset of a proposition. Since no set of descriptive states is a subset of a set of normative states, no descriptive propositions ground a normative proposition. This result, however, would be due to our semantics and our taxonomy.¹⁰ In that respect, the RAT would be subject to similar objections that modal autonomy theses face that build modal assumptions into the semantic account of normative content. The RAT would neither depend on the nature of normativity nor be sensitive to genuine counterexamples, and thus would violate G-D1 and G-D5. This illustrates that our semantics together with our assumptions about the taxonomy is not impartial. To be clear, the RAT is *not* supposed to be impartial; it makes a substantial claim that is incompatible with, e.g., naturalism. However, the semantical framework that one uses should not prejudge the truth of the RAT. This is a variant of the Impartiality Constraint. By putting emphasis on impartiality, I aim to improve on accounts according to which autonomy is a theorem of semantics.¹¹

To see that the semantical framework, as it stands, is incompatible with naturalism, consider a naturalistic view of act utilitarianism. Let D be the

proposition that act a maximizes pleasure and d the descriptive state that makes this proposition true. Let N be the proposition that act a is morally required and n the normative state that makes this proposition true. According to the view under consideration, D grounds N . However, since d is descriptive and n is normative they are distinct. Thus, it is not the case that $\{d\}$ is a subset of $\{n\}$ ($D \subseteq N|_{sN}$) and hence it is not the case that D (weakly) grounds $N|_{sN}$.

In response to this objection, I adopt a truthmaker-based semantics for ground that is sufficiently impartial. Krämer (2021) suggests the following modification of Fine's truthmaker semantics for ground:

Weak Full Ground* (\leq^*): $P_1, P_2, \dots \leq^* Q$ iff $s_1 \sqcup s_2 \sqcup \dots$ generates some truthmaker t of Q whenever s_1, s_2, \dots make P_1, P_2, \dots true, respectively.¹²

A state s generates a state t iff either s is identical to t or the proposition made true by s and only s strictly fully grounds the proposition made true by t and only t .¹³ The definition illustrates the intimate connection between generation on the level of states and grounding on the level of propositions. Generation can be described as the relationship of weak grounding on the state level. In the naturalist example, d generates n since we assumed that D grounds N and that d and n are the sole truthmakers of D and N , respectively. The modified definition allows me to accommodate views like naturalism.

One might be worried that the semantical framework is still not sufficiently impartial because it cannot account for views that assume multiple realization. Consider, for example, the normative proposition [Action a promotes the good]. Assume that this proposition can be multiply realized in descriptive ways, e.g., a promotes pleasure, a promotes knowledge, a promotes friendship, etc. If the account would require that the corresponding descriptive states are exact truthmakers of the proposition [a promotes the good], then this proposition would be descriptive, contrary to the assumption. So, in order to preserve the normative-descriptive divide on the propositional level, I shall assume that the proposition has a normative truthmaker. Formally, we can make room for this multiple-realizer view by modeling its claim as follows: the descriptive states that correspond to the descriptive realizers generate a normative truthmaker for the normative proposition [a promotes the good]. Thus, the semantical framework can accommodate the multiple-realizer view. It does not prejudge the truth of the RAT. Within this impartial framework, the autonomy thesis can be formulated on the level of states so that the result is equivalent to the RAT:¹⁴

State RAT: No descriptive state generates a normative state.

Recall that I do not primarily aim to defend the claim that normativity is autonomous; even though I am sympathetic to the idea that it is autonomous in the sense specified by the RAT. My aim is to provide an autonomy thesis that specifies an intuitively plausible and robust sense of autonomy. The RAT provides the groundwork for a comprehensive defense of the autonomy of normativity. Let me now turn to the desiderata to show that the RAT meets them and thereby improves on a contingentist account.

The RAT and the Desiderata

In this section, I show that the RAT satisfies the outlined desiderata and thereby avoids objections raised against the contingentist approach. The desideratum G-D3 deserves a detailed discussion and thus will be addressed separately in the next section.

Spurious and Genuine Counterexamples

Let me begin with the desiderata G-D2, G-D4, and G-D5. I will survey whether the RAT is immune to spurious counterexamples (G-D4), sensitive to putative genuine counterexamples (G-D5), and incompatible with particular metanormative views (G-D2).

First, I will go through some spurious counterexamples and see how the RAT can handle those cases. One can assume without loss of generality that in the disjunctive case $D = \{d\}$ and $N = \{n\}$ where d is a descriptive state and n is a normative state. $(D \vee N)|_{SN} = \{n, d \sqcup n\}$ ¹⁵ is not fully grounded in D because d does not generate any state that makes $(D \vee N)|_{SN}$ true. This is plausible since in general D is irrelevant to N .¹⁶ Hence, D does not ground $D \vee N$ in a normatively relevant way or if it does, it is a genuine counterexample which is the desired result.

In (Non-Existence), I assume that Sarah's non-existence, \bar{e} , is the sole truthmaker of [Sarah does not exist]. The normative proposition [It is not the case that Sarah ought to donate] is made true by Sarah's non-existence or by Sarah's permission to keep her money, p . If we assume that propositions are closed, i.e., that any fusion of truthmakers is a truthmaker of the proposition in question, then the proposition is also made true by the impossible state $\bar{e} \sqcup p$. Hence, [It is not the case that Sarah ought to donate]_{SN} = $\{p, \bar{e} \sqcup p\}$ is not fully grounded in [Sarah does not exist] unless \bar{e} does not generate any state that is contained in $\{p, \bar{e} \sqcup p\}$. It is plausible to assume that in this case the unless-clause is not satisfied. Hence, [Sarah does not exist] does not ground [It is not the case that Sarah ought to donate] in a normatively relevant way.

In general, whenever a proposition N has at least one descriptive and one normative truthmaker, then this proposition is (i) normative in virtue

of the normative truthmaker and (ii) if at least one of the descriptive truthmakers obtains, fully grounded by a descriptive proposition, namely the proposition that is identified with the singleton of one of N 's obtaining descriptive truthmakers. Such cases do not refute the RAT because it focuses on the normative realization of a proposition $N|_{sN}$. It requires that the grounds are relevant to the genuinely normative content of the grounded proposition.

One might ask whether the success of the RAT is restricted to disjunctive cases. Some spurious counterexamples from the literature do not seem to share the disjunctive structure. I will show that the RAT can handle those cases as well because ultimately, they have a disjunctive nature. More specifically, the grounded proposition is weakly normative. Consider a ground-based version of (Conditional Obligation):

(Ground Undertaker) That all undertakers are Church officers fully grounds that if all Church officers ought to be reverent, then all undertakers ought to be.

In order to analyze the example, one needs to identify the truthmakers of the conditional. In [Chapter 3](#), I emphasized that in general, it is not easy to tell which states make a conditional true. One possibility is to treat the grounded proposition as material conditional. However, in this case, it seems more plausible that the conditional is supposed to state a relevant connection between the antecedent and the consequent. So, which states make true the proposition [If all Church officers ought to be reverent, then all undertakers ought to be]? Recall that based on the account of truthmaker semantics for the conditional in intuitionistic logic in [Fine \(2014\)](#), we considered states that bridge the gap between the truthmakers of the antecedent and the truthmakers of the consequent to make the conditional true. The state s makes the conditional true iff every fusion of s and a truthmaker of the antecedent contains a truthmaker of the consequent. In our case, two different kinds of states satisfy this condition:

- (D-States)** States to the effect that that all undertakers are Church officers.¹⁷
- (N-States)** States that contain normative parts to the effect that some person p_i ought to be reverent.

Adding one of these states to a state that makes true the proposition [All Church officers ought to be reverent] yields a state that arguably contains a state to the effect that all undertakes ought to be reverent. Note that (D-States) are descriptive states. They also make true the proposition [If all Church officers wear green shoes, then all undertakers do]. By

contrast, (N-States) are normative states. They do not make the green-shoe-proposition true. Consider the proposition [If police officers ought to be reliable, then firemen ought to be]. Since not all firemen are in fact police officers, if this proposition is actually true, it is true in a normative way. (Ground Undertaker) is no instance of normatively relevant grounding because the normative realization of the conditional is not grounded in the descriptive proposition. The RAT handles those cases similar to the examples that were discussed before: it shows that the descriptive ground is not relevant to the normative part of the grounded proposition. Like the Re-Gap Principle, the RAT illustrates that spurious counterexamples are ultimately based on the same mechanism.

Let me discuss one last spurious counterexample.¹⁸ The descriptive proposition [Sue stands in front of children] grounds the normative proposition [Sue ought not to *flurg*]. To flurg is doing something you ought not to do in front of children. To identify the truthmakers of the grounded proposition is a complex task because of the complex structure of the proposition. However, the strategy is similar to the previously discussed cases. In order for this example to be an instance of grounding, one either must agree that the descriptive state that Sue stands in front of children is a truthmaker for the grounded proposition or that this state generates some normative truthmaker of the grounded proposition, e.g., that Sue ought not to smoke. It is unclear why a proponent of the autonomy thesis should accept the second option. Moreover, if one does, the example would at least be accompanied by (or even turn into) a genuine counterexample: Sue ought not to smoke because she stands in front of children. The first option can be handled by the RAT in the same way it handles the previous cases or more generally weakly normative propositions.

In order to show that the RAT is sensitive to putative genuine counterexamples (G-D5) and incompatible with certain metanormative views (G-D2), assume that a naturalist view is true. On a ground-based construal, this means that particular normative propositions are fully grounded in descriptive propositions. For example, assume that the following is true: that act *a* maximizes overall pleasure ($= \{d\}$) fully grounds that *a* is morally obligatory ($= \{n\}$).¹⁹ In this case, the descriptive ground is intuitively relevant to the genuinely normative parts of the grounded proposition. In order to satisfy G-D5, the RAT must be incompatible with this example.

Indeed, the RAT is incompatible with the example and thus incompatible with this form of naturalism. Note, that it is the RAT that is incompatible with naturalism rather than our semantics which is supposed to be impartial. The State RAT which is equivalent to the RAT shows this result. According to the State RAT, no descriptive state generates a normative state. However, in the naturalist case, the truthmaker of the proposition [Act *a* maximizes overall pleasure] would generate a normative state

that makes true the proposition [*a* is morally obligatory], for the former descriptive proposition strictly grounds the latter normative proposition. Therefore, a descriptive state generates a normative state. Thus, the State RAT would be false. The same result holds for the multiple-realizer case because the normative truthmaker is generated by descriptive states (the realizers). Thus, the RAT satisfies G-D5 and G-D2.²⁰

Before I turn to objections raised against Maguire's contingentist approach, let me briefly address G-D1. There is not much to say in addition to what I have argued for in the previous chapter.²¹ A claim about what grounds normative facts concerns the nature of normativity. Moreover, if normativity is autonomous in the specified sense, then the ground-based autonomy thesis can also do justice to the intuition that normativity is peculiar in virtue of that fact. At least, this is so if one thinks that many other domains are fully grounded in more fundamental domains.

Going Non-Contingentist

I shall now reconsider the objections raised against Maguire's approach. I show how the *relevant autonomy thesis* avoids the concerns and thereby generates a decisive advantage: it provides us with a substantial autonomy thesis that is sensitive to genuine counterexamples.

The Problem of Contingentism: I have argued that Maguire's account entails implausible taxonomic consequences. Recall that according to his view, the fact that it is raining or murder is morally wrong and the fact that it is raining or charity is morally wrong differ in their taxonomic status in particular situations. However, both disjunctive facts seem to be partly about rainy weather and partly about moral wrongness. Furthermore, false propositions are not subject to Maguire's autonomy thesis at all. By contrast, the RAT is not subject to this concern. The RAT is based on an essentialist account. According to *essentialism*, the taxonomic status of propositions is essential to them. The taxonomic status is determined by the truthmaker-based content of a proposition. Therefore, the RAT applies to all propositions independently of their truth-value and treats structurally similar cases in the same way.

The Problem of Triviality: I have argued that the preservation principle trivializes Maguire's autonomy thesis. The RAT avoids this concern because there are independent criteria that determine the normative and the descriptive domain. A proposition that has only descriptive full grounds might still be classified as normative like in the multiple-realizer case. As a result, the *taxonomy* is impartial with respect to metanormative views that deny the autonomy of normativity. Indeed, the RAT turns out to be false, if there is a descriptive full ground of the normative realization of a proposition.

This illustrates another important advantage of the RAT over Maguire's autonomy thesis: the RAT distinguishes between cases where the full descriptive ground of a normative proposition is relevant only to the descriptive parts of the grounded proposition and those cases where it is (additionally) relevant to its genuinely normative parts. The latter but not the former cases constitute genuine counterexamples.²² Maguire's autonomy thesis is compatible with both kinds of examples, while the RAT is just compatible with the former examples and thus satisfies G-D5. To illustrate why, I compare two examples. Consider an instance of a disjunctive case: R grounds $R \vee M$ where R is the descriptive proposition [It is raining] and M is the normative proposition [Murder is morally wrong]. Note that R is only relevant to the descriptive parts of the disjunction and irrelevant to its normative parts. Both accounts are compatible with this spurious counterexample.

Now, consider the descriptive proposition H that murder causes harm. Again, we can construct a disjunctive case: H grounds $H \vee M$. Suppose it turns out to be true that [Murder causes harm] fully grounds [Murder is morally wrong]. Accordingly, H is relevant to the genuinely normative parts of the disjunction, namely M . The RAT is incompatible with this genuine counterexample since the normative realization of $H \vee M$ is fully grounded in H by assumption. Thus, H fully grounds $H \vee M$ in a normatively relevant way. That is, if [Murder causes harm] fully grounds [Murder is morally wrong], then the RAT is false. By contrast, Maguire's autonomy thesis is compatible with the disjunctive counterexample because it is generally compatible with disjunctive cases. If there is a normative ground for $H \vee M$, then the disjunction is not just grounded in descriptive facts and if H is the sole ground, then $H \vee M$ is descriptive, according to his account. As far as I can see, there is no obvious way to achieve the intuitively correct result for Maguire's approach.

The former considerations show that the RAT is not subject to concerns raised against its predecessors. It satisfies the outlined desiderata considered so far. In a last step, I elaborate on the desideratum that I have not yet addressed, namely G-D3.

The RAT in the Metanormative Terrain

The third desideratum G-D3 requires that an autonomy thesis provides a nontrivial divide of the metanormative terrain. It already emerged that naturalist and Moorean non-naturalist views are incompatible with the RAT. However, I have not yet provided a comprehensive presentation of views that can accept the RAT. In this section, I argue that both normative realists and anti-realists alike can do so. Moreover, I show that the autonomy debate has a fruitful impact on the theoretical options that autonomist non-naturalists might explore.

Anti-Realist Positions

I already touched upon the question of whether anti-realist views can engage in the ground-based autonomy debate. I have argued that hybrid forms of non-cognitivism and quasi-realist positions can take a stand on autonomy. Once those views invoke an account of sentences like ‘Sarah’s behavior was morally praiseworthy because she saved the child’s life’, one can ask whether explanations of this form are complete. Defending a negative answer to that question is the anti-realist analogue of endorsing the RAT. It is now time to move on to another view that competes with normative realism, namely normative error theory.

Error theorists take normative sentences to purport to represent normative facts, but to be generally false in virtue of the absence of any such facts. One might think that error theorists, by design, cannot non-trivially accept the autonomy of normativity.²³ Since error theorists think that there are no *true* atomic normative sentences at all, it seems trivial for them that one cannot *get* a normative truth from a descriptive truth. However, that there are no atomic normative truths does not entail the RAT. According to the RAT, normative propositions—true or not—are not grounded in descriptive propositions in a normatively relevant way. The concept of ground that is used is non-factive. Thus, the RAT is already false if, for example, a false normative proposition is fully (non-factively) grounded in (false) descriptive propositions in a normatively relevant way. An error theorist would plausibly reject this case, and thus can accept the RAT in a non-trivial way. For if normative propositions were fully non-factively grounded in descriptive propositions, they would arguably be metaphysically continuous with descriptive propositions.²⁴ The problematic nature of normative propositions as it is assumed by error theorists would disappear. Thus, error theorists indeed should be happy to accept the RAT.

Realist Positions

Let me now turn to the realist camp. In the previous chapter, I discussed connections between an autonomy debate and the naturalism–non-naturalism debate. However, thus far, we have not considered the non-naturalist’s options in detail. Let me first explain why the devil is in the details. The RAT makes a negative claim, namely that normative propositions are not fully grounded in descriptive propositions in a normatively relevant way. This leaves us with the question of how normative propositions are grounded. Note that if there are no plausible views of how normative propositions are grounded that are compatible with the RAT, then the plausibility of the claim made by the RAT is immensely weakened. In response, I present two approaches from the literature that are compatible with the RAT.

Thereafter, I propose a third underexplored alternative approach that I take to offer a promising route in particular in light of the overall results of this book.

It already emerged that Bridge Law Non-Naturalism as defined in [Rosen \(2017\)](#) resonates with the RAT. The approach provides the following answer to the Aristotelian question: normative propositions are merely partially grounded in descriptive propositions. In addition, the descriptive partial grounds are augmented by normative bridge laws that state a connection between the relevant descriptive and normative contents. To pick up on the previous example: Sarah's behavior was morally praiseworthy in virtue of her act of life-saving and the fact that acts of life-saving are morally praiseworthy. Despite its intuitive appeal, Berker (2019) presents an influential objection against this view and its variants. He argues that normative principles are explanatorily idle. It is beyond the scope of this book to discuss Berker's objection and responses in detail. Also, to my knowledge there is nothing new that an autonomy debate could add to the controversy on this front. Accordingly, I confine myself to sketching Berker's punchline and to present another view from the literature that resonates with the RAT. Thereafter, I discuss a pertinent strand in the grounding literature and show how it provides us with an interesting novel way to respond to the Aristotelian question that is compatible with the RAT.

Let me start with the objection against the explanatory power of bridge laws. Berker distinguishes between two options: either normative bridge laws are formulated in intensional or hyperintensional terms. For concreteness, consider once more utilitarianism. An intensionally formulated principle might state that necessarily, acts are morally good *if and only if* they maximize overall happiness. The hyperintensional version states that necessarily, acts are morally good *if and only if and because* they maximize overall happiness. According to Berker, we then face a dilemma. Mere necessitation principles like the former principle do not seem to be of the right kind to get grounding claims off the ground. This can be illustrated by a rather trivial example. Necessarily, this ball is red if and only if this ball is red. However, that this ball is red together with this necessitation principle does not make it the case that this ball is red.²⁵ Thus, intensional principles do not seem to be able to close the gap between descriptive partial grounds and a normative groundee. Turning to the second option, Berker argues that there is no non-redundant explanatory work for hyperintensional normative principles. If the normative principle is true, then it is true that, for example, this act is morally good because it maximizes happiness. After all, this is what the principle tells us. According to Berker, the most natural way to interpret this claim is in terms of full ground. Thus, that the act maximizes happiness fully grounds that the act is morally good. This however renders the normative principle redundant, for it

can but need not be added to a full ground of the normative proposition, in this case that the act maximizes happiness. Thus, normative principles are explanatorily idle. This concludes my brief summary of the criticism. The Bridge Law Non-Naturalist faces a serious objection that to my knowledge still waits for a satisfying response.

Another option for the autonomist non-naturalist is to assume a distinctively normative kind of ground.²⁶ On this approach, we can and should distinguish between at least two kinds of grounding claims:

Metaphysical Ground: *P* is metaphysically grounded in Γ .

Normative Ground: *P* is normatively grounded in Γ .

For example, [The apple is red] is metaphysically grounded in [The apple is scarlet]. By contrast, [The act is morally praiseworthy] is normatively grounded in—let us say—[The act maximizes happiness]. Of course, this requires an account that settles how to distinguish between metaphysical and normative grounding and whether they are somehow related. There is a controversy about how and whether this can be done.²⁷ What matters to my current purposes is that if normative propositions are fully normatively, yet not metaphysically grounded in descriptive propositions, the normative domain would turn out to be autonomous according to the RAT.

Finally, there is an alternative approach in the grounding literature that is underexplored in the normative debate. It is on this front that an autonomy debate can add something new to the challenge for autonomist views. The general idea is that some propositions have partial grounds, yet no full ground. The notion of *merely* partial ground, i.e., a partial ground that is not part of a full ground, is recently discussed in the grounding literature.²⁸ Before I turn to the notion of merely partial ground in more detail, let me explain why it might provide non-naturalists with an attractive tool. It is nearly a truism that the normative domain depends on the descriptive domain to some extent. That Sarah's act was morally good depends on and is at least partially explained by certain descriptive features of the act of life-saving such as avoiding harm. The debate centers around the question of whether descriptive propositions *suffice* to ground normative propositions. Hence, it is nearly uncontroversial that descriptive propositions provide *partial* grounds for normative propositions. On an autonomist view, descriptive propositions do not provide full metaphysical grounds. If true, this claim entails a gap. Up to this point, I discussed two alternatives to bridge the gap. However, if Berker is correct, the choice is between two unattractive options: either the bridge principles are about grounding connections between descriptive and normative propositions, which means that they close the gap by stating that there is no gap to be closed or else

the principles are such that they are not apt to close the gap. To focus on merely partial ground is to respond to this conflict by holding on to the claim that there is a gap, but denying that it can be closed. In fact, it states that no stop-gap is required.

Let me now introduce the concept of merely partial ground in a bit more detail. I start with an example. Consider an atom *a* and some fundamental feature that atoms can possess or lack like having spin-up. We turn our attention to the proposition [Atom *a* has spin-up]. Part of what explains that atom *a* has spin-up is that atom *a* exists. However, [*a* exists] is arguably no full ground. To see why, consider a situation where *a* would have lacked the feature and instead would have possessed the feature of having spin-down. Yet, it is not at all clear which proposition could augment the proposition [*a* exists] in order to obtain a full ground for the proposition [Atom *a* has spin-up].

A second motivating case concerns completeness statements. Consider a scenario where *a* and *b* are all the things that exist. Part of what makes this proposition true is that *a* exists. What additionally contributes to the truth of this proposition is that *b* exists. Do these propositions together *fully* ground the proposition [*a* and *b* are all the things that exist]? If we assume that grounds necessitate what they ground then plausibly not. The latter proposition is incompatible with the proposition that a third object exists, while *a*'s and *b*'s existence is compatible with the existence of a third object. Even if we do not assume necessitarianism, part of what is said with the proposition [*a* and *b* are all the things that exist] is that exactly two objects exist. This, however, does not seem to be fully explained by the existence of these two objects. As a result, the propositions in question provide a partial ground, while there does not seem to be a full ground.²⁹

It is common in the literature to define partial ground in terms of full ground.³⁰ This, however, does not work if there are partial grounds without full grounds. This raises a number of ground-theoretic questions, for example, whether full ground can be defined in terms of partial ground instead. A proposal along these lines is developed in [Trogdon and Witmer \(2021\)](#). It would go beyond the scope of this inquiry to explore and discuss a general account of partial ground in more detail. For my purposes, it suffices to see that the concept of merely partial ground provides another way to be an autonomist and accept the RAT. Normative propositions are not fully grounded in descriptive propositions in a normatively relevant way—they are merely partially grounded in descriptive propositions. In other words, whenever descriptive propositions ground normative propositions in a normatively relevant way, they do so merely partially.

Let me explain why this approach seems to be attractive for non-naturalists. There is an interesting distinction between different forms of

fundamentality made in Leuenberger (2020) that essentially relies on the notion of merely partial ground. I will highlight how this distinction provides theoretically promising insights, in particular for autonomist non-naturalists. Recall that part of the Aristotelian question concerns what is derivative and what is fundamental. One objection raised against views that take particular normative truths to be fully normatively, yet not metaphysically grounded in descriptive truths is that the accounts entail that normative truths occur at the most fundamental level of reality with respect to the *metaphysical* grounding order. As a result, and as far as metaphysical fundamentality is concerned, truths about the smallest particles of the universe and—let’s say—the moral virtue of Ann’s act are on a par. To avoid this implausible result, normative truths must be moved up in the grounding hierarchy. This, however, is in tension with the denial of the claim that the normative is a mere derivative of the descriptive. The focus on merely partial ground will be instructive with respect to that tension: it is important to distinguish between *weakly* and *strongly* fundamental truths or facts. A truth P is *weakly* fundamental iff no collection of truths Γ fully ground P . In contrast, P is *strongly* fundamental iff no collection of truths Γ partially ground P .³¹ With this distinction at hand, the conflict can be addressed as follows: some physical truths, e.g., concerning smallest particles are strongly fundamental. Particular normative truths concerning moral agents or acts are not strongly fundamental because they are partially grounded in descriptive truths. This explains the observed difference between fundamental physical truths and particular normative truths. Yet, normative truths do not have full descriptive grounds. Thus, some of them are weakly fundamental. They occur on higher levels of the grounding hierarchy but are not *fully* grounded in truths belonging to the lower levels. In this sense, they are of their own kind.³²

It is worth acknowledging that this approach fits into the pattern of the overall results of this book. An autonomy thesis is often, but wrongly so, associated with claims about the plain isolation of normative truths, or claims concerning their strict independence or strict separability from the non-normative domain. These associated claims invite many of the objections pressed particularly against non-naturalism. One of the challenges concerns epistemological issues, e.g., how we could know about such normative truths. Another challenge concerns their authority or ability to guide actions given their special status.³³ I do not think that these challenges can be settled by an account along the lines I have presented here, but the notion of merely partial ground and weak fundamentality seem to provide promising candidates to mitigate the elusiveness of the normative. The normative domain is not an isolated sphere, it is related to the descriptive domain, albeit in a merely partial way which in turn explains why it is still of its own kind.

This concludes the discussion of metanormative views that can accept the RAT. The plurality of approaches that are compatible with the RAT shows that it entails a non-trivial divide of the metanormative terrain and thus satisfies the last desideratum G-D3.

Summary

In this chapter, I have developed a ground-based autonomy thesis, namely the relevant autonomy thesis. According to the RAT, the normative domain is autonomous with respect to the descriptive domain iff no normative proposition is fully grounded in descriptive propositions in a normatively relevant way. To make this intuitive claim formally precise, I proposed a truthmaker-based analysis of normatively relevant grounding according to which the normative realization of a normative proposition is not fully grounded in descriptive propositions. Within the truthmaker framework, there is an equivalent formulation of the RAT on the level of states: no normative state is generated by a descriptive state. I have argued that the RAT avoids shortcomings of existing ground-based accounts. It avoids spurious counterexamples while, at the same time, being sensitive to genuine counterexamples. It is incompatible with metanormative views that arguably deny the autonomy of normativity such as naturalism. Moreover, it entails a non-trivial divide of the metanormative terrain in general. It thereby captures a debate that concerns a wide range of metanormative views. Maybe even most importantly, the RAT makes a claim about the nature of normativity without committing its proponents to the objectionable view that the normative is isolated from the descriptive. If my arguments are correct, the RAT explicates a sense of autonomy that we should have in mind when considering the normative domain to be autonomous. I have now presented a modal and a ground-based autonomy thesis that avoid spurious counterexamples but that make a substantial claim about normativity and its relation to the descriptive domain. In the final chapter of this book, I discuss how the two autonomy theses interact and show how they thereby substantially improve our understanding of autonomy and normativity.

Notes

- 1 See, in particular, the end of [Chapter 4](#).
- 2 As I have highlighted in [Chapter 7](#), some non-naturalists disagree and argue that the claim that normative facts are fully metaphysically grounded in non-normative facts does not entail that normativity is not autonomous. For example, [Bengson et al. \(2024, p. 69\)](#) argue that it suffices for autonomy that some instantiated normative properties are deeply normative, i.e., that every set of facts encompassing their whole essence includes a normative property. I

cannot do full justice to their account here. But I'm deeply skeptical about the combination of their claims that essences that inevitably involve a normative component are *non-normative* on the one hand and that these inevitably normative essence are supposed to guarantee normative autonomy on the other hand.

- 3 See [Chapter 4](#), in particular, Section 'The Conservativeness Explication'.
- 4 I assume the *non-factive* conception of ground according to which true grounding claims do not entail that the relata are true. I will use the non-factive concept of ground throughout the rest of the book, unless otherwise indicated.
- 5 I do not use the notion of propositional part in the sense Fine uses it in Fine (2017b). My way of using it corresponds to what Fine calls *disjunctive parts*.
- 6 The autonomy thesis is a negative claim. I discuss the question of what fully grounds normative propositions in a normatively relevant way in Section 'The RAT in the Metanormative Terrain'.
- 7 Recall that the Re-Gap Principle is equivalent to the State Gap according to which no normative state exactly excludes a descriptive state.
- 8 See Fine (2012c, p. 9). Fine uses the term 'verifier' instead of 'truthmaker'. For reasons of uniformity, I changed the terminology.
- 9 See [Fine \(2012a\)](#) for a discussion of how to define strict full ground in terms of weak full ground.
- 10 [Krämer \(2021\)](#) pointed out that Fine's truthmaker-based semantics for ground rules out seemingly possible grounding structures.
- 11 Recall that, for example, the accounts in [Singer \(2015\)](#) and [Russell and Restall \(2010\)](#) entail this result.
- 12 See [Krämer \(2021, p. 48\)](#). Krämer uses the term 'verifier' instead of 'truthmaker'. For reasons of uniformity, I changed the terminology.
- 13 [Krämer \(2021, pp. 47–48\)](#). Notice that the requirement that s and t are unique truthmakers of the propositions in question is only relevant to the question of whether s generates t . Grounds and grounded proposition might have several truthmakers.
- 14 *Proof:* I show that $\neg(\text{Prop RAT}) \leftrightarrow \neg(\text{State RAT})$. Recall that the set of descriptive states \mathcal{S}^D is closed under fusion. $\neg(\text{Prop RAT}) \leftrightarrow$ Some descriptive P_1, P_2, \dots fully ground $Q|_{\mathcal{S}^N}$ for normative $Q \leftrightarrow d_1 \sqcup d_2 \sqcup \dots$ generates some truthmaker n of $Q|_{\mathcal{S}^N}$ whenever d_1, d_2, \dots make P_1, P_2, \dots true, respectively. \leftrightarrow A descriptive state, namely $d_1 \sqcup d_2 \sqcup \dots$, generates a normative state $n \leftrightarrow \neg(\text{State RAT})$.
- 15 I am assuming that disjunctions are closed under fusion.
- 16 If d generates n , then this case would be a genuine counterexample. This is the desired result, since D then would be normatively relevant to the normative realization of $D \vee N$. I will discuss this case in detail toward the end of this section.
- 17 Roughly, these states might have the following form: p_1, p_2, \dots are undertakers and Church officers and there are no other undertakers. For a detailed discussion, see Section 'Classifying Conditionals'.
- 18 Thanks to an anonymous referee for *Synthese* who raised the question of whether the RAT can handle this example.
- 19 Note that the same structure is arguably instantiated by (Promise) and (Leibniz-Euthyphro).
- 20 One might ask how the RAT relates to the examples that I discussed to motivate a restriction of the modal autonomy thesis to authoritative normativity. I shall discuss the interactions between the two autonomy theses in the next chapter, but let me briefly sketch why the situation might be different. We need

to distinguish between necessary and sufficient conditions for normative states. Only sufficient conditions are candidates for full grounds of normative propositions. Accordingly, if, for example, some normative concepts are irreducibly thick in the sense that their descriptive application conditions *fully* ground the instantiation of their normative application conditions, they plausibly also falsify the RAT. It is worth noting, however, that a relation of *partial* ground between descriptive and normative application conditions seems to suffice for classifying concepts as irreducibly thick. Proponents of the RAT thus have more room to respond.

- 21 See, in particular, Section ‘The Essential Argument’.
- 22 I discussed these cases when comparing Fine’s Gap Principle and the Re-Gap Principle.
- 23 Thanks to an anonymous referee for *Synthese* for raising this question.
- 24 See [McPherson \(2012\)](#).
- 25 For more substantial examples, see [Berker \(2019\)](#).
- 26 See, e.g., [Bader \(2017\)](#). For criticism, see [Berker \(2018\)](#).
- 27 See, e.g., [Rosen \(2017\)](#). Skeptical arguments can be found in [Berker \(2018\)](#).
- 28 See, e.g., [Dixon \(2016\)](#), [Leuenberger \(2020\)](#), and [Trogdon and Witmer \(2021\)](#).
- 29 The examples originate from [Trogdon and Witmer \(2021\)](#). Further examples are discussed in their paper. See also [Leuenberger \(2020\)](#).
- 30 See, e.g., [Fine \(2012b: 50\)](#).
- 31 See [Leuenberger \(2020\)](#), p. 2654.
- 32 I develop and defend this view in more detail in [Behrens \(2025\)](#).
- 33 For a helpful survey over some of the epistemological challenges, see, e.g., [Cuneo and Shafer-Landau \(2014\)](#). For an insightful discussion of the authority argument, see [Dasgupta \(2017\)](#).

9 The Moral of Autonomy

Taking Stock: Modal and Ground-Based Autonomy

Let me begin this chapter by summarizing the main upshots of the discussion so far. Against this background, I shall argue that the modal and the ground-theoretic strand of the autonomy debate are best treated as complementary strands. Recall that on existing logical interpretations, the autonomy thesis is in danger of becoming an overly weak claim. This result is particularly evident when purely syntactical approaches are at issue. According to these approaches, any two domains that do not overlap with respect to their characteristic vocabulary are autonomous with respect to each other. For example, the hedgehog-domain is autonomous with respect to the non-hedgehog-domain according to those accounts. Similarly, though less transparently, semantic accounts that presuppose autonomy principles, e.g., modal closure principles, to settle the taxonomic status of propositions turned out to be objectionably impartial. What is at issue then is the classification of propositions. Since autonomy is conditional on controversial classifications on these accounts, it is a rather insignificant feature. What turned out to be crucial for a *substantial* modal autonomy thesis that is immune to spurious counterexamples is the requirement that the normative content is *relevant* to the entailment relation. In order to yield a satisfying distinction between spurious and genuine counterexamples to an autonomy thesis, I have specified the relevance constraint as follows: the entailment relation holds *in virtue of* normative content; it must not hold merely in virtue of descriptive content. Based on that analysis, I have argued for a modal autonomy thesis that I have called the *Relevant Gap Principle* (Re-Gap Principle for short):

Re-Gap Principle: No descriptive premises entail a normatively relevant conclusion.

The Re-Gap Principle requires a formal account of what it means to say that an entailment relation between premises and a conclusion holds

in virtue of the conclusion's normative content. This is what I have called a normatively relevant conclusion. I have proposed a state-based analysis that employs a notion of exact exclusion between states. The result is a modal autonomy thesis that is relevantly constrained so that it resists spurious counterexamples without becoming an overly weak claim. It is an important consequence that autonomy as explicated by this account is compatible with several forms of descriptive-normative entanglements. The Re-Gap Principle is compatible with, for example, the assumption that normative assessments concerning agents presuppose their existence. However, it is an equally important feature of the principle that it is falsifiable. If it can be shown that normative content is normatively relevant to an entailment between descriptive premises and a normative conclusion, one faces a genuine counterexample to the Re-Gap Principle. I have explored one option to respond to a failure of modal autonomy as far as the normative domain as a whole is concerned. If my arguments are correct, a restriction of the Re-Gap Principle to authoritative normativity is a way to maintain a substantial modal autonomy thesis in case the general principle fails.

Relevance constraints turned out to be the key to formulate a modal autonomy thesis that can handle spurious counterexamples. Another question is whether *entailment* best captures what we have in mind when we state that we cannot *get* something normative from the non-normative. I have argued that an interpretation in terms of *ground* is better suited to accommodate that claim. The normative does not obtain *in virtue of* the exclusively descriptive. However, insights from the modal debate turned out to be relevant to the ground-based approach. Like in the modal case, the grounding relation must not hold merely in virtue of descriptive content. I proposed the following ground-based construal that I have called *Relevant Autonomy Thesis* (RAT for short):

RAT: No normative proposition is fully grounded in descriptive propositions in a normatively relevant way.

In what follows, I compare both interpretations of the autonomy thesis. It will be instructive to see the commitments they require and how their formal implementations relate.

Senses of Autonomy and Instructive Gaps

To begin with, let me sketch a pessimistic picture of the autonomy debate that one might be worried about. One might think that the intuitive plausibility of the informal autonomy thesis and likewise slogans such as No-Ought-from-Is depend on their elusiveness. Once we try to make the

claim(s) formally precise, we face a number of counterexamples. In order to immunize our interpretations against those counterexamples, the autonomy theses become complex claims that lack intuitive appeal and significance. Moreover, different interpretations concern radically distinct features some of which might not even be distinctive of normativity. This picture raises serious doubts about the significance of a normative autonomy debate. In response, let me first stress that I take the unified account of spurious counterexamples that I have defended in previous chapters and in particular the two autonomy theses themselves to provide good evidence for the inadequacy of that picture. However, more can be done. To that end, I shall first lay down what, I think, is a unified package of questions that we are interested in when engaging in the normative autonomy debate. I show that different senses of autonomy come with successively weaker commitments concerning the nature of normativity. On that picture, the RAT is the bedrock of a row of successively less committal autonomy theses. Second, I demonstrate that the formal implementation of the autonomy theses within a common framework, in particular, assumptions about the *modal*, the *mereological*, and the *priority* profile of normative states, allow us to compare the autonomy theses. The comparison sheds light on reasons to endorse or decline particular autonomy theses. The overarching aim of this chapter is to argue that the RAT and the Re-Gap Principle provide a compelling response to the outlined package of questions.

The Significance of the Autonomy Debate

I shall first state and then explain the package of questions to which a compelling account of the autonomy of normativity must respond:

- Q1: What does the claim that the normative is autonomous with respect to the non-normative mean?
- Q2: Are there different ways to understand this claim and how do they relate to each other?
- Q3: Which sort of property do the interpretations specify?
- Q4: What is normativity like if the normative domain is autonomous in the specified sense?

The discussion so far is a direct response to the first question Q1. It also comprises more or less direct answers to the remaining questions, but the current aim is to make these answers explicit. This section is mainly concerned with Q2. To be clear, I do not think that an account of the autonomy of normativity must explain how all the ways to interpret an autonomy thesis relate to each other. I have myself set aside several ways

to understand the autonomy thesis. It should be legitimate to provide one reasonable interpretation and to argue in its favor. However, the discussion of logical autonomy theses revealed that there is a tension between approaches to autonomy that take it to be a consequence of particular formal features of logic or semantics and views that insist on a genuine debate about autonomy. It seems to me that a compelling account of autonomy should be able to do some explanatory work with respect to this tension. I will discuss an example that illustrates the importance of the relation between weak and strong autonomy theses in a moment. Beforehand, let me briefly explain Q3 and Q4.

A response to Q3 helps to locate the status of an autonomy thesis among a spectrum. That is particularly important when we are concerned with substantial interpretations of the autonomy thesis. One worry is that if autonomy becomes a very strong or eccentric property of the normative (or any other) domain, then autonomy is in tension with other plausible assumptions about that domain. So, it is reasonable to clarify what kind of feature is attributed. Finally and relatedly, Q4 concerns the consequences of an autonomy debate for views of normativity. A relevant aspect here is that an autonomist position requires particular assumptions about normativity. So, part of the significance of an autonomy debate consists in its implications for a systematic account of features of normativity.

Before I address Q2 within the present framework, I want to briefly discuss an argument recently made in [Chilovi and Wodak \(2022\)](#) that is concerned with the relation between logical and metaphysical autonomy theses. I shall argue that—contrary to what the authors seem to suggest—the gap between logical and metaphysical autonomy theses that they identify is instructive. Chilovi and Wodak point out that in the pertinent literature on ethical naturalism and legal positivism, a version of Hume's Law according to which no normative proposition is *logically* entailed by a collection of descriptive propositions (henceforth *HL*) serves as premise in arguments against those views.¹ Chilovi and Wodak focus on ground-based interpretations of both naturalism and positivism. Accordingly, legal positivism is the view that for every legal fact there is a collection of social facts that fully ground the legal fact and similarly for ethical naturalism. Chilovi and Wodak argue that there is no viable route from *HL* to a denial of the ground-based theses.

The authors emphasize that there is a gap between *HL* and the denial of ground-based naturalism or positivism, respectively. The former has no direct bearing on the latter or as the authors put it: we cannot 'pull a metaphysical rabbit out of a logical hat'.² They argue that the grounding claims that are posited by positivists and naturalists, if anything, entail that grounds are *metaphysically* sufficient for the groundee (assuming necessitarianism). That, however, is compatible with the claim that in

metaphysically impossible worlds the grounds are not sufficient for the groundee. Consequently, positivists and naturalists can accept *HL* because truth-preservation between grounds and groundee fails *only* in metaphysically impossible but logically possible worlds.³ Thus, we cannot conclude from *HL* that these metanormative views are false.

While I agree with this result, it strikes me as important to be more cautious with respect to another consequence they suggest: the insignificance of *HL* outside of logic. To explain why this conclusion might be too quick, I shall do two things. First, I explain why considering modal and ground-based autonomy theses together helps to draw distinctions between spurious and genuine counterexamples. Second, I propose that *HL* is significant beyond logic because it makes us aware of reasons to endorse or deny particular autonomy theses.

Recall that a serious threat to the autonomy of normativity emerges from several putative counterexamples where it is not immediately clear what to think about them.⁴ One advantage of *HL* is that it is particularly well-suited to assess the status of what we have called spurious counterexamples. As the authors point out, *HL* is broadly acceptable even for opposing metanormative views. Even views that assume close connections between the normative and the non-normative such as legal positivism can accept *HL*. However, none of those views is committed to a particular stance toward those examples that I described as spurious. Spurious counterexamples generate complications down the line for every autonomy thesis. *HL* reveals that they are independent from assumptions about normativity. This is why they are indeed spurious.

Second, the insignificance of *HL* with respect to substantial autonomy theses reveals that it is important to clarify on what grounds one accepts *HL*. The authors seem to suggest that an autonomy thesis in terms of strict entailment is widely accepted, but irrelevant; a more substantial autonomy thesis based on metaphysical necessity is, according to them, relevant, but usually denied.⁵ They claim that views that call into question the supervenience of the normative on the non-normative (i.e., the denial of a modal autonomy thesis based on metaphysical necessity) typically do not cast doubts on whether metaphysical necessity is in play. It is important to notice that the latter claim does not seem to be correct. There is a live debate on whether to reject metaphysical supervenience and instead adopt weaker forms of supervenience.⁶ However, even if they were right about that, it seems to be interesting whether there are autonomy theses in between those views that entail a non-trivial divide of the metanormative field. Moreover, Hume's Law is often taken to have intuitive appeal. It is motivated by simple, but apparently defective examples such as the putative conclusion that an act is obligatory from the assumption that God wills it. It is not clear to me that the intuitions that these examples

prompt emerge from considering only metaphysically impossible worlds. Let me conclude the discussion with the following observation: comparing different autonomy theses is significant in order to clarify what we are aiming at with the claim that we cannot get something normative from the descriptive.

Mereological, Modal, and Priority Profile of Normative States

I now turn to the Re-Gap and the RAT and their interrelation. It turned out to be particularly useful to consider the formal implementations of the two autonomy theses within the truthmaker framework. For one thing, it allows us to reduce logical complexity on the propositional level which in turn facilitates the evaluation of both interpretations. For another thing, it makes assumptions about the normative particularly transparent. For the current purpose, the latter is particularly important. So, let me restate the state-based equivalents:

State Gap: No normative state exactly excludes a descriptive state.

State RAT: No descriptive state generates a normative state.

Whether a given model satisfies the State Gap and the State RAT depends on assumptions about normative states in the state space in three (related) respects: their *mereological* profile (P_{\sqsubseteq}), their *modal* profile (P_{\sqsupset}), and their *priority* profile (P_{\supset}). P_{\sqsubseteq} determines what is part of normative states, P_{\sqsupset} determines which states are compatible with normative states, and P_{\supset} determines which states generate normative states. One important upshot of the present account is that the very same tools of truthmaker semantics can be used to model different senses of autonomy. To see why, note that the Re-Gap Principle that is a claim about entailment relations depends on a notion of necessity that in turn leads to a distinction between consistent and inconsistent states. When I discussed the principle, I settled on a notion of conceptual necessity. However, this is not mandatory. The set of consistent states, S^{\diamond} , can represent stricter or weaker forms of necessity, e.g., logical, metaphysical, or normative necessity. Another variation in the set-up concerns the question of what states represent. For example, we might take states to be linguistic items similar to sentence letters and their negations.⁷ Roughly, different vocabulary at the linguistic level indicates different truthmakers. In order to illustrate that approach, consider the sentences ‘Xanthippe is a widow’ and ‘Xanthippe is a woman with a dead husband’. On the outlined approach, states that make the sentences true need not have anything in common.

Let me illustrate that a weak sense of modal autonomy based on such an account can be derived from minimal assumptions about normativity.

Suppose in our state space the basic states correspond to atomic sentences of a given language and their negations. The set of all states is obtained by closing the set of the basic states under fusion. It is assumed that the set of consistent states, S° , is downwards-closed, i.e., consistent states contain only consistent parts. Moreover, it is assumed that atomic sentences are incompatible with their negations, but modally independent from all other atomic sentences. Consequently, the only way in which a descriptive state can be incompatible with a normative state is by being incompatible with some of its descriptive parts. State Gap is thus true for every basic normative state because basic normative states are compatible with every descriptive state. Broadly normative states, i.e., fusions of basic normative states and descriptive states, can be incompatible with descriptive states due to their descriptive parts. This leaves us with one complication with respect to State Gap: given what has been said so far, a broadly normative state might still be wholly relevant to and thus might exactly exclude a descriptive state. On the outlined account, however, the basic normative states are modally independent from descriptive states. Thus, there is no reason to assume that any normative parts of the broadly normative state help to exclude the descriptive state. Hence, the considered state space satisfies the State Gap.

One might doubt that a state space like the one just constructed is suitable to represent normativity in any substantive sense. So, in order to make room for more substantive interpretations, one should allow for more complexity with respect to normative states as I did throughout this book. For example, I made room for a complex mereological structure of strictly normative states. Strictly normative states can have descriptive *integral* parts that cannot be separated from the normative state. The additional flexibility in the mereological profile of normative states bears on their modal profile. Strictly normative states are incompatible with states that are incompatible with their integral descriptive parts. By making room for more kinds of descriptive-normative entanglements, one poses further threats to State Gap. Consequently, one is committed to stronger assumptions about the normative in order to defend its modal autonomy. Either we reject integral descriptive parts or—as I did in the evaluation of State Gap—we argue that normative states are not wholly relevant to the exclusion of descriptive states.

At this point, State RAT can do some explanatory work. State RAT requires assumptions about the priority profile of normative states. One might constrain P_\rightarrow so that for all d and n , it is not the case that d generates n . To see why P_\rightarrow might do some work with respect to the evaluation of State Gap, suppose that a normative state n excludes a descriptive state d , while n contains a descriptive part d' that excludes d . One reason to think that n is nevertheless wholly relevant to the exclusion of d is that d' is

a maximal integral descriptive part of n that additionally generates n .⁸ This structure might be instantiated if we take some non-normative good-making feature such as *maximizing happiness* to be an essential part of *goodness*. By contrast, the existence of an agent and the agent's obligations arguably do not instantiate the same pattern. Notice that I took instances of the former to be putative genuine counterexamples to the Re-Gap Principle and instances of the latter to be merely spurious. Another aspect that deserves attention is how P_{\rightarrow} and P_{\leftarrow} relate and in particular how to interpret them. Suppose we assume that whenever d generates n , d is part of n . To be clear, one need not assume that, but if it is assumed, the assumption can be interpreted in terms of conceptual containment or in realist terms of constitution. Again, one might accept one, but deny the other.

The considerations illustrate that different views about what is at issue in the autonomy debate can be discussed within a unified framework. Moreover, I hope to have shown that the account can do substantial explanatory work on this front of the autonomy debate. It does not only make assumptions about normativity and the relevant forms of necessity formally precise, it also helps to reveal where different views come apart. To compare the success- and failure-conditions of different autonomy theses broadens our understanding of autonomy. It reveals shared assumptions as well as points of conflict.

Understanding Autonomy: Against a Dualist Picture

One central upshot of the discussed account is that the autonomy of one domain with respect to another does not entail its isolation from the other domain.⁹ This is a noteworthy result because in the debate it is often at least suggested that autonomy and isolation stand and fall together. For example, [Maguire and Woods \(2017\)](#) speak about the intuitive sense in which the ethical domain is 'insulated from the descriptive'.¹⁰ [Oddie \(2018\)](#) describes the denial of autonomy as 'the entanglement of the natural with the normative'.¹¹ According to [Jackson \(1974\)](#), if mixed terms 'cannot be analyzed into a factual element and a basic ethical element, then the autonomy of ethics is just false'.¹² The *is-ought separated model doubles* in [Schurz \(1997\)](#) can be seen as a model-theoretic manifestation of that claim.¹³

One might turn this discrepancy into an objection and claim that isolation is what we are really after. This objection concerns the third question, Q3. Let me explain why, I think, my response to Q3 is advantageous over any response that requires isolation. Serious objections that have been raised against broadly autonomist views of normativity can be seen to emerge from isolation claims. [Mackie \(1977\)](#) in his so-called *argument*

from *queerness* concludes that the denial of objective normative truths is the better option given that they are required to be of a strange sort and importantly *disconnected* from non-normative truths. More recently, [McPherson \(2012\)](#) objects to non-naturalists that their view requires unexplained necessary connections between radically distinct entities.¹⁴ The arguments concern importantly different aspects, but a basic issue seems to be that the *sui generis* status of the normative is associated with disjointness from the non-normative. This issue will come up again, when I discuss implications for normativity in the next section. For now, I focus on the notion of autonomy that the present account entails.

As I already anticipated, autonomy as explicated by the RAT and the Re-Gap Principle is importantly different. To be clear, I am not claiming that the theses provide answers to the previous objections. According to the outlined account, however, autonomy does not require isolation. Both the modal and the ground-based autonomy theses are compatible with, for example, the fact that the absence of agents entails and grounds the absence of obligations and permissions. The descriptive setup has an impact on the normative domain. It is granted that inability to act in certain ways affect the normative evaluation of those acts or that *A*'s obligation to ϕ , *A*'s permission not to ϕ , and *A*'s non-existence together constitute a mutually exclusive and with respect to a particular topic jointly exhaustive set of states. The autonomy of a particular domain Δ is compatible with entanglement between Δ and its complementary domain. Transitions from one domain into the other can be due to entanglements; what matters for autonomy is the role played by Δ -content.

The observation that a domain's autonomy need not entail its independence is also pointed out in [Fine \(2021\)](#). Recall the distinction between what he calls *Autonomy*, *Separation*, and *Independence*.¹⁵ A state space satisfies Separation if every normative state is the fusion of a descriptive and a *purely* normative component; it satisfies Independence if any possible descriptive state is compatible with any *purely* normative state. Notice that Independence and Separation presuppose the existence of purely normative states. Purely normative states, however, are particularly prone to the previously outlined objections. Fine emphasizes that Autonomy as specified by his account is weaker than and thereby advantageous over Separation and Independence. Autonomy requires that whenever a descriptive state is incompatible with a possible state, then the descriptive state is incompatible with some descriptive part of that state.

While I agree with Fine on this point, I have argued that Autonomy might be too weak. The objection might be rephrased as follows: it is an advantage of Autonomy that it does not exclude irreducible descriptive-normative entanglements; it is a disadvantage of the account that the deeper entanglements run the better things stand with respect to Autonomy. This

result is unsatisfying. The account that I have proposed instead focuses on *the way* the normative and the descriptive are entangled and thereby on different *roles* that normative content can play. Logical complexity at the propositional level generates forms of descriptive-normative entanglements so that the descriptive can bring about the normative. Key to an assessment of those cases is to reduce logical complexity by focusing on states. The focus on states reveals that it is often *just* the descriptive content that is responsible for the transition in question.¹⁶ Those and only those cases are compatible with the autonomy of normativity. The notion of autonomy that has been outlined in this book specifies a feature that makes the candidate domain *special*, but not isolated from its complement.

Partially Understanding Normativity

A related question is what normativity must be like if it is autonomous in the specified sense. This is the fourth question Q4. Let me emphasize from the outset that in this book, I was interested in an autonomy thesis that is compatible with several metanormative positions. This is why one cannot expect an autonomist position on its own to settle debates about normativity. Having said that, an autonomist position has an impact on how we construe particular notions that play a central role in normative debates. In what follows, I shall focus on the notion of *normative significance* and discuss some difficulties for autonomist realists and quasi realists that arise when normative significance is at issue.

Let me start with a realist autonomist view. According to normative realism, normative propositions are about normative facts. Autonomist realists face a tension between two features of normative facts that they wish to accommodate for. On the one hand, the claim that there is a gap between the normative and the descriptive domain requires some significant difference between the two domains. On the other hand, it seems that normativity can only play the action-guiding role that it is commonly assumed to play, if it is not disconnected from the descriptive domain. This tension raises a challenge to autonomist views: they endorse a gap that requires explanation. Which feature of the normative domain explains that gap? The main challenge is to provide an explanation that consistently combines both desired features of normative facts. As will become clear soon, strategies to explain the gap employ a notion of normative significance. However, the strategy requires that truths about *normative significance* importantly differ from truths about, for example, set-theoretical significance or the significance of determinates to their determinables (such as red and different shades of red).

In the previous section, I already anticipated how autonomy as specified by the RAT and the Re-Gap Principle helps to relieve the outlined tension.

The autonomy of normativity is compatible with the assumption that normative propositions are also about descriptive facts. Normative propositions can, for example, be partially about descriptive facts about the existence of agents. However, normative propositions are not just about descriptive facts, but still, whatever else they are about does not need to be isolated from descriptive facts. The first step for the realist autonomist to respond to the outlined challenge is to recognize that autonomy is compatible with the denial of isolation.

This step is important because it grants that descriptive facts can be normatively significant. We thereby provide the ground to be responsive to one of outlined desiderata for normative facts. However, autonomists have to make sure that there is still a gap. In order to illustrate how the autonomy debate bears on our approach to normativity, I shall consider an argument from the literature that might be used to argue for a ground-based gap. Although I cannot discuss the argument at length, I hope to show that it is central to it how we understand the employed notion of *normative relevance*. In order to avoid any confusion with relevance constraints that I have discussed earlier, I will often speak of normative significance instead. I will use the terms interchangeably in this section.

Väyrynen (2013) presents what he calls the *normative relevance argument*. According to the argument, *explanations* of normative facts in purely descriptive terms are incomplete. The notion of explanation that is used in the argument is not necessarily tied to a notion of ground.¹⁷ However, as Väyrynen points out, if it is additionally assumed that there is no explanatory gap between full grounds and groundee, the normative relevance argument entails that descriptive facts do not fully ground normative facts.¹⁸ The argument goes as follows:¹⁹ suppose a descriptive fact *D* explains a normative fact *N*. It is only if and because *D* is normatively relevant (to *N*) that it explains *N*. I will call this premise *Significance*. That *D* is normatively relevant is a fact about *D*'s normative relevance and thus a normative fact. That normative fact is not explained by *D* itself. So, an explanation of *N* just by *D* is incomplete. Given that there are no explanatory gaps between full grounds and groundee, *D* does not fully ground *N*.

One might object that even if one grants that facts about *D*'s normative significance are required to explain why *D* explains *N*, facts about *D*'s normative significance do not need to be part of a (grounding) explanation of *N*.²⁰ This criticism is indeed instructive, when one tries to better understand how an autonomist position constrains the notion of normative significance employed in the argument. Let me first turn to parallel, non-normative cases: that Socrates exists explains that singleton Socrates exists. The fact that Socrates exists can only explain the fact that singleton Socrates exists if the former is singleton-theoretically—or more generally set-theoretically—relevant. According to this reasoning, the explanatory

relation must be mediated. It would go beyond the scope of this book to enter the debate about what precisely (if anything) mediates (grounding) explanations. It has, for example, been argued that explanations have a *tripartite structure* that not only requires explanandum and explanans, but also a link between them.²¹ Others have argued that grounding explanations are always mediated by the *essences* of the items involved in the grounds or the groundee.²² What matters to the current purposes, however, is that assumptions about mediation by, for example, significance facts are commonly not taken to entail (at least in general) the incompleteness of *grounds* that do not involve those facts or to entail a non-reductionist view.²³

These considerations illustrate that it is not enough for the autonomist to claim that explanations of normative facts in descriptive terms require mediation. She must defend that these explanations are *incomplete*. For that to be plausible, facts about *D*'s normative significance must carry some explanatory weight that *D* cannot carry on its own. The work that is done by normative significance facts must exceed the work done by, for example, set-theoretical significance facts.²⁴ While it must be admitted that this is a non-trivial task for the autonomist, it also provides her with a strategy to respond to the outlined challenge: normative facts are autonomous because they possess a form of normative significance that descriptive facts cannot possess. However, normative facts can be about descriptive facts and thereby endow the descriptive facts with a derivative form of normative significance. For example, they endow them with the status of a partial ground of a normative fact.²⁵ There are arguably different ways of spelling out this general idea and it would go beyond the scope of this discussion to explore the options. However, in order to provide a better grip on the general claim, let me give one example that I take to be particularly promising: normative reasons facts. Normative reasons count in favor of or against some response such as actions or attitudes. The canonical form of normative reasons facts is [*p* is a normative reason for *A* to ϕ], where *p* is a fact or truth, *A* is some agent, and ϕ is, for example, an action. Normative reason facts endow non-normative facts with derivative normative significance when they are filled in for *p*. The non-derivative normative significance, however, is inherent in the reason relation itself and therefore in the normative reason fact.²⁶ On an autonomist position, facts that possess derivative forms of normative significance do not provide full grounds of normative facts. An autonomist position thus commits the autonomist to a demanding non-derivative notion of normative significance—at least, this is what these considerations suggest.

Let me finally consider a quasi-realist position. Although, a ground-based autonomy thesis does not require realist commitments, one should expect an autonomy thesis in quasi-realist terms to substantially differ from

a realist variant. On a quasi-realist picture, we consider the completeness of because-claims. Notice that this requires an account of because-claims in quasi-realist terms. I shall focus on an account that is developed in Berker (2020). My aim in the remainder of this section is to point out that the normative significance of descriptive propositions is also a concern for quasi-realist approaches. I shall argue that an autonomist position provides an attractive option to solve issues that arise in the context of the quasi-realist's treatment of because-claims of the form 'N because D'.

Berker (2020) points out that quasi realists need an account of because-claims like 'Kicking dogs is wrong because it causes harm' where 'because' expresses some non-causal explanatory relation like grounding. He develops a complex proposal that makes use of the notion of *basing*. I cannot discuss his proposal at length; instead, I shall briefly present the basic account and discuss how an autonomy position might help to fix a difficulty of the approach. I shall argue that the difficulty arises because the *basing* relation on its own cannot guarantee that the sentence that follows 'because' is normatively significant.

The basing relation that is discussed by Berker is supposed to be a many-one relation between mental states on the one hand and a mental state or an action on the other hand.²⁷ For example, my action of going to the shop is based on my belief that I need to buy some fruits. The quasi realist takes normative claims such as 'Kicking dogs is wrong' to express disapproval of kicking dogs. Berker suggest that the quasi realist can give an account of because-claims in terms of the basing relation. According to the proposal, a because-claim like 'Kicking dogs is wrong because it causes harm' expresses a complex state of mind in which disapproval of kicking dogs is *based* on a belief that kicking dogs causes them pain. However, Berker is aware of a complication for his proposal: we might hold an attitude like disapproval on a testimonial basis. Suppose, for example, Ann is in the mental state of disapproval of stealing based on beliefs about her mother's disapproval of stealing. It seems implausible to attribute a belief in the corresponding grounding-claim to Ann. Consequently, the account delivers the wrong verdicts in those cases. Berker suggest that we might deal with this difficulty by limiting the account to certain sorts of basing relations, e.g., non-testimonial basing.²⁸

I shall not discuss whether restrictions of that sort can withstand the initial suspicion of being *ad hoc*. Instead, I shall suggest a different diagnosis: a normative bridge element in the base might be required for it to be plausible to sincerely utter the grounding-claim. To see why, let me compare the case where Sarah is in a mental state in which her disapproval of kicking dogs is based on the belief that kicking dogs causes them pain and a case where Ann is in a mental state in which her disapproval of kicking dogs is based her mother's testimony. Why is it implausible to attribute a

belief in the corresponding grounding-claim to Ann, while it is plausible to attribute a belief in the corresponding grounding-claim to Sarah? A plausible diagnosis seems to be that Sarah's disapproval is not only based on the belief that kicking dogs causes them pain, but additionally on her disapproval of causing pain. By contrast, Ann's disapproval is only based on her mother's testimony and not also on disapproval of her mother's testimony.²⁹ Notice that the additional disapproval element in the base of the good case corresponds to a bridge principle on the realist approach like 'Causing pain is wrong'. One might object that the proposal cannot do justice to naturalist ambitions of the quasi realist's project. However, we have only employed the quasi realist's own terms to solve the difficulty. The fact that an expressivist element is indispensable for a correct analysis of normative explanations strengthens the quasi realist's claim that such an expressivist element is central to an account of normative thought and talk.

As I have flagged in the beginning of this section, further work has to be done in order to specify the details of the different autonomist positions. The RAT and the Re-Gap Principle provide the ground-work for that project. Metanormative views that wish to adopt an autonomist position can provide their specific interpretation of the formal constraints that the autonomy of normativity requires. The interpretations in turn bear on how we interpret normative notions like *normative significance* and *normative explanation*. The autonomy debate thereby contributes to a systematic account of normativity. This concludes my discussion of the response that my account provides to the package of questions raised in the autonomy debate. If my arguments are correct, the two autonomy theses provide an explanatorily powerful account of autonomy that improves our understanding of the normative domain and its relation to the descriptive domain.

Normative Sources in Scientific Contexts

In this final section, I consider questions emerging in other debates that resemble autonomy questions: questions about the source of normativity. My aim is to conclude my discussion by briefly considering other debates in which the findings of this book might be relevant. According to the principle of value-free sciences, scientific findings should not be based on non-epistemic values. The principle is challenged, however, by the observation that scientists work with normative tools like thick normative concepts and normative models. Debates about thick normative concepts in science directly relate to what has been discussed in this book. So, in what follows, I will only briefly outline the central concern. The debate about the role of normative models in science is less directly related. Normative models

give rise to the question of whether and how models can exert *normative guidance*. These questions concern sources of normativity, in particular regarding normative models that are used in *artificial intelligence systems*. I shall consider putative limits of normative models and possible constraints on the data basis that result from taking into account autonomy.³⁰

Let me start with thick normative concepts. Recently, canonical examples of thick normative concepts are amplified by more inconspicuous cases such as disease, creativity, or colleague.³¹ Moreover, in social sciences, putatively thick normative concepts like well-being, biodiversity, vulnerability are ubiquitous. It has been argued that these concepts present a challenge for science because in general science is not equipped with an expertise on normative questions.³² The question arises how research that involves such concepts can be done in responsible ways. The weight of this challenge crucially depends on the outcome of autonomy debates. Suppose the non-normative suffices to entail or explain (ground-theoretically) the normative. The challenge would be limited because researchers could focus on the non-normative basis. There might remain an epistemic challenge that is concerned with the question of how to get to know which non-normative features are normatively relevant, but it would in principle be legitimate to employ those concepts on a scientific basis. The assessment of the usage of thick normative concepts in science crucially depends on our stance toward autonomy.

In the remainder of this section, I take a look at the recently growing debate about normative models in science and their adequacy conditions.³³ I take this debate to be particularly interesting because the role normative models play seems to be similar to the roles played by normative arguments and normative explanations: we have a non-normative input and a normative output, but it is unclear what else, if anything, is needed for the transition to work. I start by briefly introducing normative models. I point out that it is unclear what must be built into normative models for them to be able to guide action. Questions that arise in that debate are to some extent parallel to questions that I discussed in the context of the autonomy debate. It will be instructive to see whether insights gained in the previous discussion can add something to this debate as well.

A distinction can be made between descriptive and normative models, or more precisely between different roles played by a model.³⁴ A model that plays a descriptive role is supposed to represent a particular part of the world to facilitate predictions, explanations, and to discover patterns. A model that plays a normative role aims to provide normative guidance to agents.³⁵ Models that play a normative role are salient in formal epistemology, decision theory, economics, and the development of AI. I shall be concerned with the question of how normative models can successfully exert normative guidance toward agents. Beck and Jahn (2021) point out

that there is little systematic discussion of that question. I will follow them in calling the property of normative models in virtue of which they exert normative guidance the *normative source* of that model.

It is important to acknowledge that the authors aim to explain a demanding notion of normative guidance. To see that, two partial functions of normative models that are assumed by the authors are particularly important. On the one hand, a normative model recommends a certain course of action, on the other hand, it provides a reason for an agent to act in the recommended way. The second partial function distinguishes normative models from mere heuristics that direct agents to normative verdicts that are independently justified. According to Beck and Jahn, normative models are supposed to provide *additional* reasons for the verdicts they generate. So, the model itself does normative work.³⁶ One of the examples that the authors discuss are *expected utility theory models* in a context of medical decisions. Medical professionals are supposed to act in their client's interests.³⁷ According to the authors, these models are built in three steps. First, agents are presented with hypothetical choices that they have to respond to. Second, the agents are modeled based on descriptively adequate models of decision-making. Lastly, an expected utility theory model is derived from the descriptively adequate model that is taken to be adequate to determine how the agents *should* act.³⁸ The normative model might, for example, recommend that the medical professional stops life-sustaining measures. Recall that the normative model is supposed to provide an *additional* reason for the agent, in that case the medical professional, to act in the recommended way. That the model recommends stopping life-sustaining measures provides a reason to do so.

Beck and Jahn argue that the normative source of normative models consists of two parts. On the one hand, they assume an intrinsic feature of normative models to reproduce a *projectable pattern* and on the other hand, they assume an array of independently justified normative verdicts. The idea is that the model extends normative justification to cases of uncertainty.³⁹ The normative model exerts guidance to agents partially in virtue of being a concise, pattern-revealing summary of the independently justified normative verdicts and partially due to a projection of the pattern onto novel situations. A first thing to notice is that the autonomy of normativity is compatible with the account. According to this approach, independently justified normative verdicts are part of the normative source of the model. So, one might think that the way that normative models work is similar to normative bridge principles in arguments and in normative explanations. We add normative content that comprises normative patterns to non-normative content to get normative verdicts. What seems particularly interesting about Beck and Jahn's account, however, is that

they assume that models increase the number of justified verdicts over the number of such verdicts that can be obtained by model-independent sources.

A number of interesting questions in the context of autonomy considerations arise. Let me start with a toy-example. The authors assume that the transitivity assumption in expected utility theory models according to which we should prefer *A* over *C*, if we prefer *A* over *B*, and *B* over *C* is such a normative pattern. Suppose we have model-independent reason to prefer bananas over apples in situations in which we prefer bananas over pears, and pears over apples. In addition, we have model-independent reason to prefer becoming a philosopher over becoming a pop star in situations in which we prefer becoming a philosopher over becoming a mathematician, and becoming a mathematician over becoming a pop star. Our data instantiates a pattern. In order to obtain a projectable pattern, we have to generalize features of the data. To that end, we must assume that novel cases resemble the data sufficiently to be subsumable under the same pattern. Note that similarity in one or the other respect is cheap. The central question is what guarantees that the model identifies normatively significant similarities.

A general feature of models might shed some light on the discussion. As Titelbaum (forthcoming) emphasizes, a model always requires an interpretation. An interpretation of a model includes a *domain of applicability* of the model. The importance of this point can be illustrated by appeal to the AGM approach of belief revision.⁴⁰ The models intended domain of applicability is restricted to particular situations, namely situations in which an agent learns by gaining certainties that are to be retained in future learning experiences. Outside the domain the model entails implausible (normative) results. So, what seems crucial to the success of normative models is that they are applied to cases where the extending function preserves the normative justification that is provided by the data basis. For that to be the case, the normative model must be such that it identifies the *right* patterns (whatever they are). It is unclear how to understand ‘the right pattern’, if normative models are supposed to be more than mere heuristics. I am skeptical about the claim that normative models have an intrinsic feature of generating projectable patterns that extend justification. This suggests that the normative work is done (if it can be done) by the normative data basis and the interpretation of the model. The data can only provide a basis for projection, if it instantiates a sufficiently robust pattern.⁴¹ This raises the important question of how the normative data must look like in order to provide an adequate basis. If normative principles must be built into normative models, then what we gain by normative models is limited. A number of questions deserve further attention, but cannot be discussed here. I hope to have illustrated,

though, that we should care for the sources of normativity and how we might get normative verdicts from normative models.

Summary

In this chapter, I have argued that the ground-based and the modal strand of the autonomy debate make complementary contributions to a better understanding of the autonomy of normativity. Autonomy can be interpreted so that it comes with close to none commitments concerning the nature of normativity; other interpretations require substantial commitments. Gaps between those interpretations are instructive because they help to clarify on what grounds we take the normative domain to be autonomous. My preferred autonomy theses, the RAT and the Re-Gap Principle, and their implementation within the truthmaker framework provide a unified way to deal with spurious counterexamples while they can be discussed and compared in a common formal set-up. Assumptions about the profile of normative states make the relations between different autonomy theses particularly transparent. It is a strength of the account that it does not entail the isolation of the normative domain. Moreover, its flexibility allows different metanormative views to interpret the formal constraints in a way that respects their specific view on normativity. I have argued that autonomy questions shed light on normative notions such as normative significance and bear on how we employ normative models and concepts in science.

The intuitive appeal of the claim that we cannot get something normative from the non-normative has been threatened by spurious counterexamples that refute overly simple interpretations of the autonomy thesis. The availability of intuitively plausible interpretations that resist spurious counterexamples and explain why they are indeed spurious reveals the relevance of autonomy theses in normative debates.

Notes

- 1 See Chilovi and Wodak (2022) and references therein.
- 2 Chilovi and Wodak (2022, p. 638).
- 3 The authors phrase their argument in terms of absolute necessity. Since the argument does not depend on that particular notion, I will not say anything more specific about it here. It suffices that there is a form of necessity that is stricter than metaphysical necessity.
- 4 The canonical examples are the ones in [Prior \(1960\)](#).
- 5 Chilovi and Wodak (2022, p. 461).
- 6 For example, Hattiangadi (2018) argues that non-naturalists ought to reject metaphysical supervenience. Rosen (2020) provides an argument against metaphysical supervenience that assumes non-naturalism as a premise. In response, he suggests a supervenience claim in terms of normative necessity.

- 7 See Fine (2021, p. 892, fn. 3), for the proposal to model syntactic approaches to autonomy in that way.
- 8 This might provide us with a formal account of *merely partial ground*. A state t that is part of another state s but that neither can be fused with other proper parts of s to obtain s nor generates s is a merely partial ground. The equivalence between unsupplemented parts within the truthmaker framework and merely partial grounds is discussed in Leuenberger (2020). A detailed discussion has to be left for another occasion.
- 9 This observation is also made in Fine (2021). Autonomy as specified by his account does not require what he calls Independence or Separation. I turn to his distinction toward the end of this section.
- 10 Maguire and Woods (2017, p. 424).
- 11 Oddie (2018, p. 607).
- 12 Jackson (1974, p. 94).
- 13 See Schurz (1997, pp. 97–98).
- 14 Another interesting argument along similar lines is given in Dasgupta (2017). The argument raises doubts on whether *sui generis* properties as posited by some autonomist views can be shown to play the role they want them to play. The main issue seems to be that it is unclear how *sui generis* properties manage to be action guiding given that they are not suitably related to non-normative facts about us.
- 15 Fine (2021, p. 901). I have discussed Fine’s account in detail in Chapter 4.
- 16 Notice that I thereby take some steps toward Separation, though in a harmless way. I do not presuppose that there are purely normative states, but I require that we can determine whether the normative state adds something in comparison to its descriptive part.
- 17 In fact, Väyrynen uses the argument to motivate doubts about the project of understanding normative explanations in terms of ground. See Väyrynen (2013, p. 165).
- 18 This assumption has been made, for example, in Fine (2012b, p. 39). See Schaffer (2017) for the claim that explanatory gaps are wide-spread.
- 19 I mainly follow the presentation in Väyrynen (2013, pp. 158–159).
- 20 Väyrynen (2013, p. 160) discusses this objection, but argues that D ’s insufficiency to explain N is supported by the observation that repeating ‘Because D ’ is not responsive to someone’s query concerning N .
- 21 See, e.g., Glazier (2016).
- 22 See, e.g., Fine (2002) and Rosen (2010). For a defense of the view that the grounds ground the grounding facts, see, e.g., Bennett (2011).
- 23 For discussion in the metanormative context, see, e.g., Heathwood (2012). See, e.g., Dasgupta (2014) in the context of physicalism.
- 24 Väyrynen suggests that the justificatory role of normative explanations and their special relation to reasons for action is a factor that makes normative explanations disanalogous to other grounding explanations (Väyrynen 2013, p. 173). In previous chapters, I have discussed an autonomist position according to which some normative facts are weakly fundamental in the sense that they are partially grounded in non-normative facts, but do not have any full grounds. I am sympathetic to the claim that reason-facts are weakly fundamental in that sense. See also Endnote 30 in this chapter.
- 25 Väyrynen distinguishes between two explanatory questions one of which concerns *bearers* of normativity while the other concerns *sources* of normativity (Väyrynen 2013, p. 171).

- 26 For this reason, it is promising to analyze right- and good-makers, i.e., facts about an entity that make it right or good, in terms of normative reasons. I develop this proposal in more detail in Behrens (ms).
- 27 See Berker (2020, pp. 10–11).
- 28 Berker (2020, p. 18, fn. 29).
- 29 Note that this proposal might require that there are basic forms of disapproval that are not based on disapproval of something else or cases of mutual basing. I lack the space to discuss these issues here, but note that the idea that there is mutual grounding on the most fundamental level is considered in the grounding literature. See, e.g., Giannotti (2021).
- 30 Another field where the findings of the discussion might turn out to be fruitful is the application of a theory of weak fundamentality in the reasons-debate. It is commonly assumed that the reason relation applies to facts. Consequently, it seems plausible to assume that reason-facts are partially grounded in the non-normative fact the reason relation applies to. Reason-first views assume that reason-facts are strongly fundamental within the normative domain. A notion of weak fundamentality seems particularly interesting to those views. Reason-facts could be weakly fundamental with respect to the whole grounding hierarchy. See Behrens (2025).
- 31 For recent work on thick normative concepts in social sciences, see van der Weele (2021) and references therein. Herfeld and Djordjevic (2021) discuss thick concepts in economics. In the context of aesthetics, Sánchez-Dorado (2020) provides an argument for the claim that creativity is a thick epistemic concept.
- 32 See Alexandrova and Fabian (2021). I read this as a claim about some, but not all branches of science.
- 33 See, e.g., Colyvan (2013), Beck and Jahn (2021), and Titelbaum (forthcoming).
- 34 See, e.g., Buchak (2013). Notice that some models play both roles or different roles in different contexts. Some approaches reject distinctively normative models and instead take them to be descriptive models of an *ideal* agent such as the perfect Bayesian updater.
- 35 I am here following Beck and Jahn's characterization of normative models (Beck and Jahn 2021, p. 125). Titelbaum (forthcoming) introduces a normative model as a model that attempts to fit normative facts, while descriptive models attempt to fit descriptive facts. Notice that Beck and Jahn's notion is more demanding. I shall grant that some models are supposed to be normative models in their sense.
- 36 See Beck and Jahn (2021, p. 141).
- 37 See Bleichrodt et al. (2001).
- 38 See Beck and Jahn (2021, p. 130).
- 39 See Beck and Jahn (2021, pp. 141–144).
- 40 See Alchourrón et al. (1985).
- 41 See Queloz (2025) for elaborated doubts concerning the systematicity of truth in the normative domain.

10 Conclusion

This concludes my defense of an account of what it is for the normative domain to be autonomous with respect to the non-normative domain. I began this book with an alleged truism that in its most straightforward interpretations is demonstrably false: we cannot get something normative from the non-normative. I have defended a primary ground-based and a complementary modal interpretation of that claim. According to the ground-based thesis, RAT, no normative proposition is fully grounded in non-normative propositions in a normatively relevant way. According to the modal thesis, the Re-Gap Principle, no non-normative premises entail a normatively relevant conclusion. Both interpretations do justice to the intuitive appeal of the initial claim and resist spurious counterexamples that prove simple interpretations of the claim to be false.

Central to my account is that it entails a *unified explanatory* approach to the cluster of examples that pose a threat to the autonomy of normativity. Logically complex propositions give rise to two sources of spurious counterexamples to autonomy theses. One source is the false classification of propositions. The distinction between normative and non-normative propositions requires a fine-grained, hyperintensional account of content. The truthmaker-based account reveals that some apparently non-normative propositions have normative content. What is more, normative propositions can be true in non-normative ways. It is crucial to the impartiality of the account that the employed taxonomy does not assume taxonomic principles that presuppose forms of autonomy. The account thus improves on existing approaches that incur such a commitment. The precise analysis of normative content is relevant for debates beyond the autonomy debate.

Another source of spurious counterexamples is indifference with respect to the role that is played by the normative content. Transitions from the non-normative to the normative do not refute the autonomy of normativity, if the normative content is irrelevant to the transition in question. Only if non-normative propositions ground (or entail) a normative proposition *in virtue of* the normative content, do we face a genuine counterexample

to the ground-based (or modal) autonomy thesis. The key assumption of my account is that relevance constraints play a central role in specifying a plausible sense of autonomy.

If the arguments I gave are correct, an autonomy thesis in terms of ground best represents what we have in mind when we consider the claim that we cannot get something normative from the non-normative. The literature on the autonomy of normativity, though, focuses on modal interpretations. The novel perspective reveals that we are not only concerned with whether, but *why* normative propositions are true. I have argued that this question not only matters for metanormative views that are classically engaged in the debate like proponents of realism, but concerns the majority of the metanormative field.

Naturally, some issues could only be addressed briefly and some important questions that follow on a satisfying account of autonomy such as a comprehensive defense of the autonomy of normativity could not be fully approached within the limits of this book. I hope to have shown, though, that the account I have defended is a rich and powerful account. The formal implementation of the autonomy theses within the truthmaker framework provides the groundwork for a comprehensive defense by specifying the assumptions it would require. The flexibility of the formal account respects specific approaches to normativity. This feature makes a fruitful comparison of different versions of the autonomy thesis possible. I hope to return to strands that I had to leave at some point in the discussion in future research. In particular, the truthmaker framework in its application in normative contexts, the notion of weak fundamentality, and the notion of normative significance can do further important work in normative and metanormative debates.

In closing, let me explain why I emphasized the need to specify a falsifiable autonomy thesis. Recall [Pigden's \(1989\)](#) skeptical question: 'what reason do we have to believe it [the autonomy thesis] true, besides the, perhaps temporary, drying up of counterexamples?'. At least to me, it seems that a compelling reason to believe that normativity is autonomous depends on what exactly normativity is. This question, however, still waits for a satisfying response. What I hope to have shown, though, is that we can dry up spurious counterexamples. That provides the ground for an important debate about genuine cases and shows that an autonomy thesis is relevant.

Bibliography

- Alchourrón, C. E., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2), 510–530. <https://doi.org/10.2307/2274239>
- Alexandrova, A., & Fabian, M. (2021). Democratising measurement: Or why thick concepts call for coproduction. *European Journal for Philosophy of Science*, 12(1), 1–23. <https://doi.org/10.1007/s13194-021-00437-7>
- Angell, R. B. (1989). Deducibility, entailment and analytic containment. In J. Norman & R. Sylvan (Eds.), *Directions in relevant logic* (pp. 119–143). Kluwer Academic Publishers.
- Asay, J. (2013). Truthmaking, metaethics, and creeping minimalism. *Philosophical Studies*, 163(1), 213–232. <https://doi.org/10.1007/s11098-011-9808-0>
- Ayer, A. J. (1936). *Language, truth and logic*. V. Gollancz.
- Bader, R. M. (2017). The grounding argument against non-reductive moral realism, *Oxford Studies in Metaethics*, 12, 106–134. <https://doi.org/10.1093/oso/9780198805076.003.0005>
- Baker, D. C. (2017). The varieties of normativity. In T. McPherson & D. Plunkett (Eds.), *The Routledge handbook of metaethics* (pp. 567–581). Routledge.
- Baker, D. C. (2018). Skepticism about ought simpliciter, *Oxford Studies in Metaethics*, 13, 230–252. <https://doi.org/10.1093/oso/9780198823841.003.0011>
- Beck, L., & Jahn, M. (2021). Normative models and their success. *Philosophy of the Social Sciences*, 51(2), 123–150. <https://doi.org/10.1177/0048393120970908>
- Behrens, S. (2025). Don't mind the gap: how non-naturalists should explain normative facts. *Philosophical Studies*, 182, 1221–1242. <https://doi.org/10.1007/s11098-025-02312-0>
- Behrens, S. (ms). No Guide to Ground: Right-making and Right-makers.
- Behrens, S. (2021). No normative free lunch: Relevance and the autonomy of the normative domain. *Synthese*, 199(5), 13163–13186. <https://doi.org/10.1007/s11229-021-03371-6>
- Behrens, S. (2024). A semantics for moral error theory. *Analysis*, 84(2), 221–230. <https://doi.org/10.1093/analys/anad047>
- Bengson, J., Cuneo, T., & Shafer-Landau, R. (2023). The source of normativity. *Mind*, 132(527), 706–729. <https://doi.org/10.1093/mind/fzac063>
- Bengson, J., Cuneo, T., & Shafer-Landau, R. (2024). *The moral universe*. Oxford University Press. <https://doi.org/10.1093/9780198914594.001.0001>
- Bennett, K. (2011). By our bootstraps. *Philosophical Perspectives*, 25(1), 27–41. <https://doi.org/10.1111/j.1520-8583.2011.00207.x>
- Berker, S. (2018). The unity of grounding. *Mind*, 127(507), 729–777. <https://doi.org/10.1093/mind/fzw069>

- Berker, S. (2019). The explanatory ambitions of moral principles. *Noûs*, 53(4), 904–936. <https://doi.org/10.1111/nous.12246>
- Berker, S. (2020). Quasi-dependence. *Oxford Studies in Metaethics*, 15, 195–218. <https://doi.org/10.1093/oso/9780198859512.003.0009>
- Berto, F., & Jago, M. (2018). Impossible worlds. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University.
- Bleichrodt, H., Pinto, J. L., & Wakker, P. (2001). Making descriptive use of Prospect theory to improve the prescriptive use of expected utility. *Management Science*, 47(11), 1498–1514. <https://doi.org/10.1287/mnsc.47.11.1498.10248>
- Bolzano, B. (1837). *Wissenschaftslehre*, Bd. 1–4. Sulzbach: Seidelsche Buchhandlung.
- Brown, C. (2014). Minding the is-ought gap. *Journal of Philosophical Logic*, 43(1), 53–69. <https://doi.org/10.1007/s10992-012-9253-3>
- Brown, C. (2015). Two versions of Hume's law. *Journal of Ethics and Social Philosophy* (1), 2–7. <https://doi.org/10.26556/JESP.V9I1.170>
- Buchak, L. (2013). *Risk and rationality*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199672165.001.0001>
- Chilovi, S., & Wodak, D. (2022). On the (in)significance of Hume's law. *Philosophical Studies*, 179(2), 633–653. <https://doi.org/10.1007/s11098-021-01674-5>
- Clarke-Doane, J. (2014). Moral epistemology: The mathematics analogy. *Noûs*, 48(2), 238–255. <https://doi.org/10.1111/j.1468-0068.2012.00875.x>
- Clipsham, P. (2015). Developing a post-prior taxonomy of ethical sentences. *Philosophia*, 43(3), 801–820. <https://doi.org/10.1007/s11406-015-9602-x>
- Colyvan, M. (2013). Idealisations in normative models. *Synthese*, 190(8), 1337–1350. <https://doi.org/10.1007/s11229-012-0166-z>
- Correia, F. (2010). Grounding and truth-functions. *Logique Et Analyse*, 53(211), 251–279.
- Cullity, G., & Gaut, B. (1997). *Ethics and practical reason*. Oxford University Press. <https://doi.org/10.1111/1467-9213.00160>
- Cuneo, T., & Shafer-Landau, R. (2014). The moral fixed points: New directions for moral nonnaturalism. *Philosophical Studies*, 171(3), 399–443. <https://doi.org/10.1007/s11098-013-0277-5>
- Dasgupta, S. (2014). The possibility of physicalism. *Journal of Philosophy*, 111(9–10), 557–592. <https://doi.org/10.5840/jphil20141119/1037>
- Dasgupta, S. (2017). XV—Normative non-naturalism and the problem of authority. *Proceedings of the Aristotelian Society*, 117(3), 297–319. <https://doi.org/10.1093/arisoc/aox016>
- Dixon, T. (2016). Grounding and supplementation. *Erkenntnis*, 81(2), 375–389. <https://doi.org/10.1007/s10670-015-9744-z>
- Dreier, J. (2002). Meta-ethics and normative commitment. *Philosophical Issues*, 12(1), 241–263. <https://doi.org/10.1111/j.1758-2237.2002.tb00069.x>
- Edgington, D. (2020). Indicative conditionals. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University.
- Enoch, D. (2010). How objectivity matters. *Oxford Studies in Metaethics*, 5, 111–152.
- Fantl, J. (2006). Is metaethics morally neutral? *Pacific Philosophical Quarterly*, 87(1), 24–44. <https://doi.org/10.1111/j.1468-0114.2006.00246.x>
- Faroldi, F. (2014). Denial of responsibility and normative negation. In F. Cariani, D. Grossi, J. Meheus, & X. Parent (Eds.), *Deontic logic and normative systems*, Springer. https://doi.org/10.1007/978-3-319-08615-6_7
- Faroldi, F. (2019). *Hyperintensionality and normativity*. Springer Verlag. <https://doi.org/10.1007/978-3-030-03487-0>

- Fine, K. (2002). Varieties of necessity. In T. S. Gendler & J. Hawthorne (Eds.), *Conceivability and possibility* (pp. 253–281). Oxford University Press.
- Fine, K. (2005). *Modality and tense: Philosophical papers*. Oxford University Press. <https://doi.org/10.1093/0199278709.001.0001>
- Fine, K. (2010). Some puzzles of ground. *Notre Dame Journal of Formal Logic*, 51(1), 97–118. <https://doi.org/10.1215/00294527-2010-007>
- Fine, K. (2012a). Counterfactuals without possible worlds. *Journal of Philosophy*, 109(3), 221–246. <https://doi.org/10.5840/jphil201210938>
- Fine, K. (2012b). Guide to ground. In F. Correia, & B. Schnieder (Eds.), *Metaphysical grounding* (pp. 37–80). Cambridge University Press.
- Fine, K. (2012c). The pure logic of ground. *Review of Symbolic Logic*, 5(1), 1–25. <https://doi.org/10.1017/s1755020311000086>
- Fine, K. (2014). Truth-maker semantics for intuitionistic logic. *Journal of Philosophical Logic*, 43(2–3), 549–577. <https://doi.org/10.1007/s10992-013-9281-7>
- Fine, K. (2016). Angelic content. *Journal of Philosophical Logic*, 45(2), 199–226. <https://doi.org/10.1007/s10992-015-9371-9>
- Fine, K. (2017a). A theory of truthmaker content I: Conjunction, disjunction and negation. *Journal of Philosophical Logic*, 46(6), 625–674. <https://doi.org/10.1007/s10992-016-9413-y>
- Fine, K. (2017b). A theory of truthmaker content II: Subject-matter, common content, remainder and ground. *Journal of Philosophical Logic*, 46(6), 675–702. <https://doi.org/10.1007/s10992-016-9419-5>
- Fine, K. (2017c). Truthmaker semantics. In B. Hale, C. Wright, & A. Miller (Eds.), *A companion to the philosophy of language* (pp. 556–577). John Wiley and Sons Ltd. <https://doi.org/10.1002/9781118972090.ch22>
- Fine, K. (2021). Truthmaking and the Is-Ought Gap. *Synthese*, 198(2), 887–914. <https://doi.org/10.1007/s11229-018-01996-8>
- Finlay, S. (2006). The reasons that matter. *Australasian Journal of Philosophy*, 84(1), 1–20. <https://doi.org/10.1080/00048400600571661>
- Finlay, S. (2010). Recent work on normativity. *Analysis*, 70(2), 331–346. <https://doi.org/10.1093/analys/anq002>
- FitzPatrick, W. J. (2008). Robust ethical realism, non-naturalism, and normativity. *Oxford Studies in Metaethics*, 3, 159–205. <https://doi.org/10.1093/oso/9780199542062.003.0008>
- Foot, P. (1958). Moral arguments. *Mind*, 67(268), 502–513. <https://doi.org/10.1093/mind/LXVII.268.502>
- Foot, P. (1972). Morality as a system of hypothetical imperatives. *The Philosophical Review*, 81(3), 305–316. <https://doi.org/10.2307/2184328>
- Fraassen, B. C. V. (1969). Facts and tautological entailments. *Journal of Philosophy*, 66(15), 477–487. <https://doi.org/10.2307/2024563>
- Geach, P. T. (1965). Assertion. *Philosophical Review*, 74(4), 449–465. <https://doi.org/10.2307/2183123>
- Gemes, K. (1994). A new theory of content I: Basic content. *Journal of Philosophical Logic*, 23(6), 595–620. <https://doi.org/10.1007/BF01052779>
- Gemes, K. (1997). A new theory of content II: Model theory and some alternatives. *Journal of Philosophical Logic*, 26(4), 449–476. <https://doi.org/10.1023/A:1004291217792>
- Giannotti, J. (2021). Fundamental yet grounded. *Theoria*, 87(3), 578–599. <https://doi.org/10.1111/theo.12293>
- Gibbard, A. (2003). *Thinking how to live*. Harvard University Press.

- Glazier, M. (2016). Laws and the completeness of the fundamental. In M. Jago (Ed.), *Reality making* (pp. 11–37). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198755722.003.0002>
- Hansson, S. O. (2014). Deontic diversity. In F. Cariani, D. Grossi, J. Meheus, & X. Parent (Eds.), *DEON 2014* (pp. 5–18). Springer International Publishing Switzerland. https://doi.org/10.1007/978-3-319-08615-6_2
- Hare, R. M. (1952). *The language of morals*. Oxford University Press. <https://doi.org/10.1093/0198810776.001.0001>
- Hare, R. M. (1964). The promising game. *Revue Internationale de Philosophie*, 18(70), 398–412. https://doi.org/10.1007/978-1-349-15336-7_15
- Hattiangadi, A. (2018). Moral supervenience. *Canadian Journal of Philosophy*, 48(3–4), 592–615. <https://doi.org/10.1080/00455091.2018.1436034>
- Heathwood, C. (2012). Could morality have a source? *Journal of Ethics and Social Philosophy*, 6(2), 1–19. <https://doi.org/10.26556/jesp.v6i2.62>
- Herfeld, C., & Djordjevic, C. (2021). Thick concepts in economics: The case of Becker and Murphy’s theory of rational addiction. *Philosophy of the Social Sciences*, 51(4), 371–399. <https://doi.org/10.1177/004839312111008541>
- Hesse, M. (1970). Theories and the transitivity of confirmation. *Philosophy of Science*, 37(1), 50–63. <https://doi.org/10.1086/288279>
- Hill, S. (2008). ‘Is’-‘Ought’ derivations and ethical taxonomies. *Philosophia*, 36(4), 545–566. <https://doi.org/10.1007/s11406-008-9131-y>
- Horwich, P. (2018). Is truth a normative concept? *Synthese*, 195(3), 1127–1138. <https://doi.org/10.1007/s11229-016-1208-8>
- Howard, N. R., & Laskowski, N. G. (forthcoming). Robust vs. formal normativity II, or: No gods, no masters, no authoritative normativity. In D. Copp, & C. Rosati (Eds.), *The Oxford handbook of metaethics*. Oxford University Press.
- Humberstone, I. L. (1982). First steps in a philosophical taxonomy. *Canadian Journal of Philosophy*, 12(3), 476–478. <https://doi.org/10.1080/00455091.1982.10716342>
- Humberstone, I. L. (1996). A study in philosophical taxonomy. *Philosophical Studies*, 83(2), 121–169. <https://doi.org/10.1007/bf00354286>
- Humberstone, L. (2019). Recent thought on is and ought: Connections, confluences, and rediscoveries. *Journal of Applied Logics*, 6, 1373–1446.
- Jackson, F. (1974). Defining the autonomy of ethics. *Philosophical Review*, 83(1), 88–96. <https://doi.org/10.2307/2183875>
- Jackson, F. (1991). Decision-theoretic consequentialism and the nearest and dearest objection. *Ethics*, 101(3), 461–482. <https://doi.org/10.1086/293312>
- Jackson, F. (1998). *From metaphysics to ethics: A defence of conceptual analysis*. Oxford University Press. <https://doi.org/10.1093/0198250614.001.0001>
- Jackson, F. (2013). Autonomy of ethics. In H. LaFollette (Ed.), *The international encyclopedia of ethics* (pp. 459–465). Blackwell Publishing. <https://doi.org/10.1002/9781444367072.wbiee015>
- Jago, M. (2020). Truthmaker semantics for relevant logic. *Journal of Philosophical Logic*, 49, 681–702. <https://doi.org/10.1007/s10992-019-09533-9>
- Jago, M. (2023). Conjunctive and disjunctive parts. In F. Faroldi, & F. V. D. Putte (Eds.), *Outstanding contributions to logic: Kit fine* (pp. 167–188). Springer. https://doi.org/10.1007/978-3-031-29415-0_9
- Karmo, T. (1988). Some valid (but no sound) arguments trivially span the ‘Is’-‘Ought’ Gap. *Mind*, 97(386), 252–257. <https://doi.org/10.1093/mind/XCVII.386.252>
- Kiesewetter, B. (2011). ‘Ought’ and the perspective of the agent. *Journal of Ethics and Social Philosophy*, 5(3), 1–24. <https://doi.org/10.26556/jesp.v5i3.57>

- Kim, D. (2024). Exact truthmaker semantics for modal logics. *Journal of Philosophical Logic*, 53(3), 789–829. <https://doi.org/10.1007/s10992-024-09752-9>
- Kirchin, S. (2013). II—Evaluation, normativity and grounding. *Aristotelian Society Supplementary Volume*, 87(1), 179–198. <https://doi.org/10.1111/j.1467-8349.2013.00225.x>
- Krämer, S. (2017). A hyperintensional criterion of irrelevance. *Synthese*, 194(8), 2917–2930. <https://doi.org/10.1007/s11229-016-1078-0>
- Krämer, S. (2021). Singular troubles with singleton Socrates. *Philosophy and Phenomenological Research*, 103(1), 40–56. <https://doi.org/10.1111/phpr.12692>
- Krämer, S., & Roski, S. (2017). Difference-making grounds. *Philosophical Studies*, 174(5), 1191–1215. <https://doi.org/10.1007/s11098-016-0749-5>
- Lange, M. (2018). What would normative necessity be? *Journal of Philosophy*, 115(4), 169–186. <https://doi.org/10.5840/jphil2018115412>
- Leary, S. (2021). What is normative non-naturalism? *Ergo*, 8(52), 787–814. <https://doi.org/10.3998/ergo.2253>
- Leibniz, G. W. (1989). Reflections on the common concept of justice. In L. E. Loemker (Ed.), *Philosophical papers and letters* (pp. 561–573). Springer Netherlands. https://doi.org/10.1007/978-94-010-1426-7_60
- Leuenberger, S. (2020). The fundamental: Ungrounded or all-grounding? *Philosophical Studies*, 177(9), 2647–2669. <https://doi.org/10.1007/s11098-019-01332-x>
- Lewis, D. (1988). Relevant implication. *Theoria*, 54(3), 161–174. <https://doi.org/10.1111/j.1755-2567.1988.tb00716.x>
- Mackie, J. L. (1977). *Ethics: Inventing right and wrong*. Penguin Books. <https://doi.org/10.2307/2184791>
- Maguire, B. (2015). Grounding the autonomy of ethics. *Oxford Studies in Metaethics*, 10, 188–215. <https://doi.org/10.1093/acprof:oso/9780198738695.003.0008>
- Maguire, B. (2017). The autonomy of ethics. In T. McPherson, & D. Plunkett (Eds.), *The Routledge handbook of metaethics* (pp. 431–442). Routledge.
- Maguire, B., & Woods, J. (2017). Model theory, Hume’s dictum, and the priority of ethical theory. *Ergo*, 4, 419–440. <https://doi.org/10.3998/ergo.12405314.0004.014>
- Maguire, B., & Woods, J. (2020). The game of belief. *Philosophical Review*, 129(2), 211–249. <https://doi.org/10.1215/00318108-8012843>
- Maitzen, S. (1998). Closing the ‘Is’-‘Ought’ gap. *Canadian Journal of Philosophy*, 28(3), 349–366. <https://doi.org/10.1080/00455091.1998.10715977>
- Maitzen, S. (2010). moral conclusions from non-moral premises. In C. R. Pigden (Ed.), *Hume on Is and Ought* (pp. 290–309). Palgrave MacMillan.
- Maurin, A. (2019). Grounding and metaphysical explanation: It’s complicated. *Philosophical Studies*, 176(6), 1573–1594. <https://doi.org/10.1007/s11098-018-1080-0>
- Mavrodes, G. (1968). On deriving the normative from the non-normative. *Papers of the Michigan Academy of Arts and Sciences*, 53, 353–365.
- McPherson, T. (2012). Ethical non-naturalism and the metaphysics of supervenience. In R. Shafer-Landau (Ed.), *Oxford Studies in metaethics* (Vol. 7, pp. 205–234). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199653492.003.0007>
- McPherson, T. (2018). Authoritatively normative concepts. In R. Shafer-Landau (Ed.), *Oxford Studies in metaethics* (Vol. 13, pp. 253–277). Oxford University Press. <https://doi.org/10.1093/oso/9780198823841.003.0012>
- McPherson, T. (2021). Supervenience in ethics. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University.

- McPherson, T., & Plunkett, D. (2020). Conceptual ethics and the methodology of normative inquiry. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *conceptual engineering and conceptual ethics* (pp. 274–303). Oxford University Press. <https://doi.org/10.1093/oso/9780198801856.003.0014>
- McPherson, T., & Plunkett, D. (2024). The fragmentation of authoritative normativity. In R. Shafer-Landau (Ed.), *Oxford Studies of metaethics* (Vol. 19, pp. 1–28). Oxford University Press. <https://doi.org/10.1093/oso/9780198911876.003.0001>
- Moore, G. E. (1903). *Principia ethica*. Dover Publications. <https://doi.org/10.2307/2176289>
- Nelson, M. T. (2007). More bad news for the logical autonomy of ethics. *Canadian Journal of Philosophy*, 37(2), 203–216. <https://doi.org/10.1353/cjp.2007.0018>
- Nietzsche, F. (1998). *On the genealogy of morality*. Cambridge University Press.
- Oddie, G. (2018). Non-naturalist moral realism, autonomy and entanglement. *Topoi*, 37(4), 607–620. <https://doi.org/10.1007/s11245-017-9454-z>
- Parfit, D. (2011). *On what matters: Two-volume set*. Oxford University Press. <https://doi.org/10.1093/acprof:osobl/9780199572816.001.0001>
- Pigden, C. (2010a). *Hume on Is and Ought*. Palgrave-Macmillan.
- Pigden, C. (2010b). Substance, content, taxonomy and consequence: A comment on Stephen Maitzen. In C. Pigden (Ed.), *Hume on Is and Ought* (pp. 313–319). Palgrave-Macmillan.
- Pigden, C. R. (1989). Logic and the autonomy of ethics. *Australasian Journal of Philosophy*, 67(2), 127–151. <https://doi.org/10.1080/00048408912343731>
- Plunkett, D. (2020). Conceptual truths, evolution, and reliability about authoritative normativity. *Jurisprudence*, 11(2), 169–212. <https://doi.org/10.1080/20403313.2020.1715104>
- Plunkett, D., & Sundell, T. (2013). Disagreement and the semantics of normative and evaluative terms. *Philosophers' Imprint*, 13(23), 1–37.
- Popper, K. (1962). *Conjectures and refutations: The growth of scientific knowledge*. Routledge. <https://doi.org/10.2307/2218271>
- Prior, A. (1960). The autonomy of ethics. *Australasian Journal of Philosophy*, 38(3), 199–206. <https://doi.org/10.1080/00048406085200221>
- Prior, A. N. (1949). *Logic and the basis of ethics*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198241577.001.0001>
- Queloz, M. (2025). Can AI rely on the systematicity of truth? The challenge of modelling normative domains. *Philosophy and Technology*, 38(34), 1–27. <https://doi.org/10.1007/s13347-025-00864-x>
- Roberts, D. (2011). Shapelessness and the thick. *Ethics*, 121(3), 489–520. <https://doi.org/10.1086/658898>
- Rosen, G. (2010). Metaphysical dependence: Grounding and reduction. In B. Hale, & A. Hoffmann (Eds.), *Modality: Metaphysics, logic, and epistemology* (pp. 109–135). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199565818.003.0007>
- Rosen, G. (2017). Metaphysical relations in metaethics. In T. McPherson, & D. Plunkett (Eds.), *The Routledge handbook of metaethics* (pp. 151–169). Routledge.
- Rosen, G. (2020). What is normative necessity? In M. Dumitru (Ed.), *Metaphysics, meaning, and modality*, Oxford University Press. <https://doi.org/10.1093/oso/9780199652624.003.0011>
- Rusnock, P., & George, R. (2014). *Bernard Bolzano: Theory of science*. Oxford University Press.

- Russell, G. (2021). How to prove Hume's law. *Journal of Philosophical Logic*, 51(3), 603–632. <https://doi.org/10.1007/s10992-021-09643-3>
- Russell, G., & Restall, G. (2010). Barriers to implication. In C. R. Pigden (Ed.), *Hume on Is and Ought* (pp. 243–259). Palgrave MacMillan.
- Rynin, D. (1957). The autonomy of morals. *Mind*, 66(263), 308–317. <https://doi.org/10.1093/mind/LXVI.263.308>
- Sánchez-Dorado, J. (2020). Novel & worthy: Creativity as a thick epistemic concept. *European Journal for Philosophy of Science*, 10(3), 1–23. <https://doi.org/10.1007/s13194-020-00303-y>
- Scanlon, T. (1998). *What we owe to each other*. Harvard University Press. <https://doi.org/10.2307/j.ctv134vmrn>
- Schaffer, J. (2000). Trumping preemption. *Journal of Philosophy*, 97(4), 165. <https://doi.org/10.2307/2678388>
- Schaffer, J. (2009). On what grounds what. In D. Manley, D. J. Chalmers, & R. Wasserman (Eds.), *Metametaphysics: New essays on the foundations of ontology* (pp. 347–383). Oxford University Press. <https://doi.org/10.1093/oso/9780199546046.003.0012>
- Schaffer, J. (2012). Grounding, transitivity, and contrastivity. In F. Correia & B. Schnieder (Eds.), *Metaphysical grounding: Understanding the structure of reality* (pp. 122–138). Cambridge University Press.
- Schaffer, J. (2016). Grounding in the image of causation. *Philosophical Studies*, 173(1), 49–100. <https://doi.org/10.1007/s11098-014-0438-1>
- Schaffer, J. (2017). The ground between the gaps. *Philosophers' Imprint*, 1–26.
- Schnieder, B. (2006). Truth-making without truth-makers. *Synthese*, 152(1), 21–46. <https://doi.org/10.1007/s11229-004-7905-8>
- Schnieder, B. (2020). Grounding and dependence. *Synthese*, 197(1), 95–124. <https://doi.org/10.1007/s11229-017-1378-z>
- Schroeder, M. (2009). Hybrid expressivism: Virtues and vices. *Ethics*, 119(2), 257–309. <https://doi.org/10.1086/597019>
- Schurz, G. (1991). Relevant deduction. *Erkenntnis*, 35(1-3), 391–437. <https://doi.org/10.1007/BF00388295>
- Schurz, G. (1997). *The is-ought problem an investigation in philosophical logic* (Vol. 1 of *Trends in Logic*). Kluwer Academic Publishers.
- Schurz, G. (2010). Comments on Restall, Russell and Vranas. In C. Pigden (Ed.), *Hume on Is and Ought* (pp. 268–271). Palgrave-Macmillan.
- Searle, J. R. (1964). How to derive 'Ought' from 'Is'. *The Philosophical Review*, 73(1), 43–58. <https://doi.org/10.2307/2183201>
- Shorter, J. (2010). Professor prior on the autonomy of ethics. In C. Pigden (Ed.), *Hume on Is and Ought* (pp. 47–48). Palgrave MacMillan.
- Singer, D. J. (2015). Mind the Is-Ought Gap. *Journal of Philosophy*, 112(4), 193–210. <https://doi.org/10.5840/jphil2015112412>
- Strevens, M. (2008). *Depth: An account of scientific explanation*. Harvard University Press. <https://doi.org/10.2307/j.ctv1dv0tnw>
- Titelbaum, M. G. (forthcoming). Normative modeling. In J. Horvath (Ed.), *Methods in analytic philosophy: A contemporary reader*. Bloomsbury Academic Press.
- Trogon, K., & Witmer, D. G. (2021). Full and partial grounding. *Journal of the American Philosophical Association*, 7(2), 252–271. <https://doi.org/10.1017/apa.2020.26>
- Urmson, J. O. (1953). Some questions concerning validity. *Revue Internationale de Philosophie*, 7(3), 217–229.

- Van der Weele, S. (2021). Thick concepts in social research: What, why, and how? *International Journal of Qualitative Methods*, 20. <https://doi.org/10.1177/16094069211066165>
- van Roojen, M. (2018). Moral cognitivism vs. non-cognitivism. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University. <https://doi.org/10.1177/16094069211066165>
- Väyrynen, P. (2013). *The lewd, the rude and the nasty: A study of thick concepts in ethics*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199314751.001.0001>
- Väyrynen, P. (2021). Against moral contingentism. *Thought: A Journal of Philosophy*, 10(3), 209–217. <https://doi.org/10.1002/tht3.494>
- Väyrynen, P. (2021). Thick ethical concepts. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University.
- Vranas, P. (2010). Comments on ‘Barriers to Implication’. In C. Pigden (Ed.), *Hume on Is and Ought* (pp. 260–267). Palgrave-Macmillan.
- Wedgwood, R. (2004). The Metaethicists’ mistake. *Philosophical Perspectives*, 18(1), 405–426. <https://doi.org/10.1111/j.1520-8583.2004.00035.x>
- Wolf, A. (2020). Ruling out solutions to Prior’s dilemma for Hume’s law. *Thought: A Journal of Philosophy*, 9(2), 84–93. <https://doi.org/10.1002/tht3.445>
- Yablo, S. (2014). *Aboutness*. Princeton University Press. <https://doi.org/10.1515/9781400845989>
- Yablo, S. (2025). Relevance without minimality. In P. van Elswyk, D. Kindermann, C. D. Kirk-Giannini, & A. Egan (Eds.), *Unstructured content*. Oxford University Press.

Index

- \wedge/\vee -closure 15–16, -desideratum 45
- act utilitarianism 175, 182
- action guiding 215; and action-guiding role 206
- actual* oughts 148, 150–51
- AGM 213
- Alchourròn, C. E. 216
- alethic-deontic: logics 109; and principles 34
- Alexandrova, A. 216
- all-things-considered 135–6, 154–5
- analytic entailment 83–6, 99, 125
- Angell, R. B. 99
- anti-realists 7, 54, 162, 188–9; accounts 171
- argument from queerness* 204
- Aristotelian: questions 190, 193; and *View* 167
- artificial intelligence systems 211; AI 211
- Asay, J. 24
- authoritative: BETTER 135–7, 141–6, 151, 156; -/non-authoritative distinction 156; normative 134; normative states 147–50; normativity 99, 125, 133–6, 139–52, 198; OUGHT 134; REASON 135–7, 155; REQUIREMENT 135
- authoritativeness 134–5, 138
- authority 133, 144; and authority-revealing concepts 137
- autonomist: non-naturalist 169, 191; position 63, 66, 68, 75, 104, 179–80, 191–2, 204, 206, 210
- autonomists 63, 66, 72, 88
- autonomy: of normativity 1, 3, 88, 91, 123, 129, 131, 166, 184, 187, 189, 199, 212, 214, 217–18; principles 173, 197; and theses 1, 4–7
- Baker, D. C. 140–4, 146
- basic excluders 113
- Beck, L. 211–12
- Berker, S. 170, 190, 209
- black holes 106
- bridge law 168, 190; and Bridge Law Non-Naturalism 168, 180, 190–1
- bridge principles 63–4, 66, 109, 191, 210, 212
- Brown, C. 41, 43
- categoricity 135
- causation 145
- Chilovi, S. 200
- closure: principle 19, 29, 37, 69, 172, 197; under negation 18–19
- completeness statements 192
- composition 17–18
- conceptual truths 34, 101, 132, 137
- conjunction 15–17, 53, 56–9, 61, 84, 129, 131; elimination 18
- conjunctive part 59, 61, 118–19
- conjunctive-disjunctive-part (*cd*-part) 118–19
- connected 109
- Conservativeness Explication, The 81, 83, 85–6, 98, 107, 151; *see also* conservativeness principle
- conservativeness principle, the 81–5; refined 82–4
- consistency 15–16; and -desideratum 27, 41, 173

- constraints 15, 113, 139, 211;
 arbitrariness 148; conceptual 34,
 45; descriptive 139, 152; formal
 210, 214; impartiality 15, 20, 23,
 182; metaphysical 38; neutrality
 20; non-trivial descriptive 139;
 relevance 4, 98, 111, 121, 198, 207,
 218; substance 137
- content containment 84
- contingent: descriptive disjunct 28, 36;
 facts 42, 150; sentence 106
- Contingentism 171, 176, 184
- converse entailment 20, 23, 36
- Converse Metaphysical Autonomy
 (CMA) 172, 174, 176
- correctness 15–16; and -desideratum 56
- decisiveness 145
- deductive theory confirmation 106
- deflationary: mathematical pluralism
 140; pluralism 140, 146; pluralist
 134, 150
- deontic: accessibility relation 3; logic
 30–1; operators 30, 32, 108;
 tautologies 33
- descriptive: components 47, 55,
 92–4, 96–7, 122–3, 129; content
 16, 26, 29, 44–5, 58, 75, 92,
 101, 110, 114, 119, 122, 129,
 197; facts 8, 164, 169–70,
 172–5, 180, 188, 207–8; *see also*
 DESCRIPTIVITY-content
- descriptive-normative: distinction 34,
 53–4; entanglements 125, 128–9,
 132, 152, 198, 203
- DESCRIPTIVITY-content 125, 128–9,
 132, 152, 198, 203
- difference-maker 22, 39, 107; account
 109; difference-making 22, 107
- disjunction 3, 15–18, 24, 34, 36, 40–1,
 44–5, 52–4, 56–9, 71, 73–5, 88;
 mixed 28, 87–8; -introduction 18
- distinctively 134, 179
- domain 7, 17, 85, 166; of applicability
 213; of discourse 21–2
- double-indexed semantic framework
 35
- double-negation introduction 18
- Downwards-N-Closure 55, 66; and
 Downwards-S^o-Closure 111
- downwards-closed 203
- Dreier, J. 21, 41
- drug case 130
- eccentric property 200
- embedding problem 170
- Enoch, D. 22
- equivalence 103
- error theory 150; normative 189;
 and error theorists 189
- essence-based approaches 167
- essentialism 40, 187; and essentialist
 view 74
- Ethical Naturalism 168, 200
- Ethics: Inventing Right and Wrong* 161
- exact: exclusion 110–15, 124, 198;
 incompatibility 110; truthmakers
 51, 53, 71, 111–12, 118, 183;
 truthmaking 51, 111
- exactly exclude 6, 111–18, 120–5,
 128–30, 146, 149, 152, 166, 202–3
- exclusion strategy 34
- expected utility theory models 212–13
- explanations: complete 171;
 incomplete 208; normative 210–2;
 pragmatic 62
- explanatory: basis 23, 29; and thesis
 152
- expressivism 170
- extension-fragility 31–32
- extension-revealing 137–8
- extrinsic feature 15
- factive: operators 60, 72, 75;
 propositions 72, 75
- falsemaker 53, 56, 61, 66, 70, 96, 121
- Fantl, J. 21
- Faroldi, F. 62
- Finean: account 94, 123, 101;
 approach 101, 122–3, 182; Gap
 Principle 98, 118, 120, 121–4, 152;
 and Fine's truthmaker semantics
 182–3
- Fine, K. 5–6, 51, 53–4, 64–5, 91–2,
 94–6, 185, 205
- Foot, P. 135
- formal 3, 6, 13, 17–8, 20, 24, 28, 33,
 35, 39, 51, 53, 55, 58, 74–5, 105–6,
 112, 114, 118, 120, 134–6, 144,
 146–9, 180–3, 194, 197–200, 202,
 204, 210–1, 214, 218; and informal
 16, 22, 33, 39, 83–4, 96–7, 113, 198

- fragility 30–1, 34; *see also* extension-fragility; translation-fragility
- Frege-Geach problem 170
- fully ground: 4, 167, 180–90, 192–4, 198, 200, 207, 217; -ed 3, 161, 168–9, 174–6, 181–2, 184–9, 192–4, 198, 217
- fundamental 1, 139, 161–2, 164, 167, 187, 192, 193, 218
- fusion 44, 52–3, 55, 57, 59, 61, 65–8, 71, 73, 94, 103, 112–3, 115–6, 121, 184–5, 203, 205, 207
- Gap Principle 91–2, 94–5, 97–8
- General Relevant Gap Principle 120
- God's eye perspective 130
- ground-based: autonomy 8, 176; autonomy thesis 2, 4, 6, 8, 162, 163, 166–8, 170–1, 173–4, 176, 179–81, 187, 194, 205; debate 162, 168–9, 189
- ground-based explication: of autonomy 163, 166; of naturalism 169
- Ground Undertaker 185
- ground-theoretic: questions 192; and sense 168
- grounding 14, 161–3, 167, 172–6, 181–183, 186, 193–4, 209; claims 162, 174, 182, 190–1, 200, 209–10; disjunctions 172; explanation 207–8; metaphysical 193; principles 167; relation 1, 2, 162, 198; truth 167
- GUN 66
- HALLOW 164–6; -example 17; -propositions 165–6
- Hansson, S. O. 136
- hedgehog-domain 83–6, 116, 197
- heuristics 212–13
- Howard, N. R. 135
- Humberstone, I. L. 17–19, 37–8; and Humberstone-Lewis-results 17–18
- Hume's Law* 1, 200–1
- hybrid 44, 170, 189
- hyperintensional: account 5, 69, 190, 217; distinctions 46–7, 49, 75, 148; notion 40, 190; semantics 47
- ideal worlds 30–2, 54, 87, 89
- if-then-statements 170
- incompatibility 93–5, 110, 112, 114, 122–3, 149; requirements 114
- incompleteness 208
- inconsistent 52, 103, 113, 117; and state 52, 111, 113, 202
- incremental conditional 64–6, 70–1
- independence 91, 94, 103, 205; strict 193
- inexact exclusion 111, 114, 124
- inexactness 111
- inference-relative-vacuity 82
- integral descriptive parts 55–6, 96, 105, 119, 203
- integral variant 138,
- intensional: account 4–5, 26, 44–5, 69, 75, 90; and taxonomies 43
- internal feature 142
- intrinsically 34
- invariantist semantics 37
- irreducibly thick 128–9
- irrelevance spectrum 102, 104–5, 120–1
- irrelevant: conclusions 102, 118, 120, 124; completely 107–8, 117; normatively 119; partially 98, 107–10, 114–6, 120
- Is-Ought Gap 133
- Is-Ought inferences 109
- is-ought separate model doubles 204
- Jackson, F. 130, 204
- Jahn, M. 211–12
- Karmo, T. 19, 41
- knowledge 5; normative 74–5
- Krämer, S. 118, 183
- Laskowski, N. 135
- legal positivism 200–1
- Leibniz-Euthyphro 6
- Leuenberger, S. 193
- Lewis, D. 17
- Lewisian subject-matter 38
- logical autonomy 28–29, 83, 97, 101, 152; debate 163, 171, 176, 180; theses 5, 8, 79, 81, 83, 85–8, 91–4, 98, 105, 107, 109, 120–1, 128, 151, 163–4, 171, 200; and simple logical autonomy theses 3, 19, 24, 28–29, 69, 97, 113, 120, 163
- logically necessary truths 69–70, 72

- MA 174, 176
Mackie, J. L. 162, 204
Macron 102, 109, 111–13, 161, 163–4
Maguire, B. 20, 34, 36–7, 171–5, 204
Maguire's autonomy thesis 171–6, 188;
and contingentist approach 187–8
material conditional 64–5, 71, 185
McPherson, T. 134–5, 137, 139, 205
merely partial ground 191–3
mereological profile 123, 125, 148–9,
202–3
metanormative: debate 7, 133, 218;
positions 21, 206; views 7, 14, 20,
22, 169–170, 176, 180, 184, 186–7,
194, 201, 210, 214, 218
metaphysical autonomy 172
Metaphysical Autonomy Theses 5,
7–8, 13–15, 71, 152, 159, 175,
200; simple 163
metaphysical ground 191, 193
metaphysical supervenience 201
minimal truth-value 109
minimalism 14
minimalist: account 54; position 14;
view 14
minimality requirement 112
mixed: conjunctions 44, 57–8, 61, 65;
disjunctions 28, 34, 36, 40, 42,
44–5, 54, 57–9, 61, 63, 65, 81,
87–8, 171, 173, 180
modal autonomy thesis 4, 6, 8, 105,
110, 114, 118, 120, 124–5, 136,
197–8, 201; substantial 197–8
modal profile 102–105, 124, 128, 146,
148–51, 202–3
modalized state space 52, 103–4
model-theoretic taxonomies 20, 33
Moorean 8; connections 169; non-
naturalist views 188; open-question
arguments 137; views 169
moral: laws 15–16, 23, 68, 71–2,
75; oughts 130, 142, 145, 150;
propositions 15; rationalism 148–9;
rule 139
moral-reason-giving 136
morally wrong 26–7, 40–2, 45, 68, 94,
132, 170, 173–4, 187–8
multiple-realizer view 183
MURDER 129
murder 40–2, 45, 68, 129, 173–4,
187–8
N-making feature 169
naturalism 168–9, 174–5, 182–3,
186, 194, 200; –non-naturalism
distinction 169
naturalist 168–9, 175, 180, 186, 201,
210
Naturalist views 168–9, 180
nature: of authoritativeness 138; of
normativity 1, 3, 167, 174, 179,
182, 187, 194, 199, 214
NC 26–7
necessary truth 34, 44–5, 53, 60,
69–70, 72
necessitarianism 192, 200
negation 17–22, 24, 28, 33, 36, 53,
60–5, 68, 75, 98, 202–3; external
62, 65; internal 62
neutrality 20–4; Neutrality_{NC} 21;
Neutrality_{ND} 22
No-Ought-from-Is 198
NOC 28–9
non-arbitrariness 139
non-arbitrary 136, 139–40, 151
non-autonomist 19, 88, 91, 95–6, 166,
169, 179
non-cognitivism 189
non-cognitivist views 13, 21
non-contributing 111, 115, 122
non-difference-making: approach 107;
and occurrence 121
non-ethical 19
non-existence 61, 96, 122–3, 149, 184,
205
non-factive: grounding 174; notion of
ground 163; operators 74
non-monotonic 113
non-naturalism 168–9, 193
non-naturalist 168–9, 180, 188, 191,
205
non-normative claim 2; good-making
204; truth 29, 188, 205
non-obtaining fact 174
non-paradigmatic 40, 172–3, 178
non-perspectivism 129–30
non-reductive naturalism 175
non-symmetric 113
non-trivial 19–20, 56, 103, 105, 136,
149–50; divide 180, 194, 201;
normative content 44–6; and task
17–18, 208
Non-Triviality Constraint 17–19

- norm invariance 35, 89
norm-variant 35–7, 39, 89
normative autonomy 8, 13, 30, 35, 63, 82–3, 85, 98, 124, 131, 134, 151–2, 164, 168, 179–80
normative: components 26; content 2, 4, 6, 8, 16, 27–9, 42, 44–7, 53–5, 57–9, 70–1, 75, 86, 93, 96–8, 101–2, 104–5, 110, 114, 117, 120, 122–4, 128, 131–2, 152, 181–2, 185, 190, 197–8, 206, 212, 217; contribution 123, 149; deflationary pluralism 134, 140, 150; difference-makers 39; explanations 210–11; expression 26, 54; extension 31; facts 8, 54, 162, 164, 168–9, 172–5, 179, 187, 206–8; falsemaker 56, 121; ground 190–1; guidance 211, 212; irrelevance 116; negation 60–1, 63; nihilism 22; nihilist 22, 28, 60–2, 68; propositions 4, 16–17, 19, 21, 27, 29–31, 33–4, 36–40, 47, 54, 56–8, 60–3, 66, 69, 71, 74–5, 86–7, 89, 91, 97, 128–9, 131–4, 136, 145–8, 151–2, 164, 168, 176, 181, 186, 189–92, 206–7, 217–18; realism 206; realization 181–2, 185–8, 194; relevance 98, 133, 142, 145–7, 149–52, 181, 20; relevance argument 207; significance 133, 148, 206–10, 214, 218; source 144, 210, 212; standard 35; states: broadly 55–6, 203; purely 56, 61, 94, 119, 205; strictly 56, 203; status 41, 74, 149, 173, 181; structure 30, 33–4, 4; translation 31; truthmakers 54, 56–8, 61, 63, 66, 68–71, 75, 84, 92, 104, 117–18, 181, 183–7; truths 29, 145, 148, 161, 166, 169–70, 189, 193, 205
normative-belief propositions 165
normative-descriptive divide 183
normative-knowledge propositions 165
normatively irrelevant 89, 114, 116–20, 128
normatively relevant grounding 181–2, 186, 194
null-state 52, 55, 113
O-irrelevant 108–9; -operator 108; -restricted 108
objective 4, 137, 139–40, 143, 145–7, 150, 161, 205; and normative relevance 145
obligation 7, 45, 55, 71, 96, 119, 123, 125, 133, 139, 143, 146, 205
Oddie, G. 43–6, 57
Old Testament 15
ontological commitment 14, 27, 29
orthodoxy 15–16; and -desideratum 68, 71
Ought-implies-Can Principle 60
ought-operator 31, 108–9
Ought-statement 1
overdetermination 97, 109–10, 115
Overgeneralization Problem, The 95
overriding 144
Parfit, D. 135
partial ground 163, 190–3, 208
partial irrelevance 101, 115
particular propositions 87
permission state 124
perspectivism 129–30
Pigden, C. 82–3
Pigden's argument-relativism 40–1
Plato 14, 52
Plunkett, D. 134–5, 137
POLITE 137
possible-worlds 26, 28–9, 35, 38–9, 43, 46–7, 51–2, 69, 71, 95, 111, 149–50, 201; and semantics 35, 92
practical ought 134
practical relevance 145
preservation principle 175, 187
Prior-style 81
Prior, A. N. 3, 69, 82
Prior's examples 3–4, 19, 45, 82
Problem of Contingentism, The 173, 187
Problem of Triviality, The 187
projectable pattern 212–3
promise 45, 123–5, 133; and promissory-rules 133
Prop RAT 182
propositional parts 52, 58, 118–19; sub- 26–7
prudential: ought 134, 136, 144, 148; and rule 139
psychological or social relevance 150

- quasi-realist 189, 208–9
Quinean View 167
- RAT, The 188–92, 194, 199, 202, 205–6, 210, 214
- Re-Gap 120, 124, 202; approach 122–4; *Principle* 98–9, 121–5, 128–9, 131–3, 136, 146, 148–9, 151–2, 166, 186, 197–9, 202, 205–6, 210, 214, 217
- Re-Gap's exception 131
- realist 7, 54, 168–70, 189, 204, 206–10; and autonomist view 206
- realized content 58–61, 74–5
- reason-implying 135
- reductivist tendency 180
- refined conservativeness principle 83–4; Refined Conservativeness_{ANA} 84; and Refined Conservativeness_{VAC} 82
- relativity theory 106
- relativism: argument- 40–1; speaker- 40–3; taxonomic 40, 41; world- 40–2
- relevance: -based account 141, 143, 145; (inter-)subjective 146; sociological 146; and spectrum 114, 122
- Relevant Autonomy Thesis 179, 181, 187, 194, 198
- Relevant Gap Principle 120, 197
- residuation state 65–6
- Restall, G. 30–4, 87
- restricted: autonomy thesis 131; inference barriers 86, 88–91; and Re-Gap Principle 151–2, 166
- ROC 29
- rodent-concepts 28
- Rosen, G. 168, 180, 190
- rule-implying 135
- Russell, G. 30–34, 87, 89
- Russell's proposal 89–90
- salva validitate* 107–9
- Schaffer, J. 167
- Schroeder, M. 170
- Schurz, G. 105–9, 204
- Schurz's syntactic account 107, 109
- scientific theories 106
- Searle, J.R. 133
- Searle's argument 133
- semantic account 4, 34, 107, 110, 147, 182, 197
- sentential connectives 162
- separability 85, 94, 98, 205
- separable 55–6, 81, 83, 85–6, 179
- separation 92, 94
- set of states 53, 112, 115
- set-theoretical significance 206, 208
- SHMACTFUL 138–9
- significance: 3, 17, 19, 63, 65, 83, 145, 166, 170, 199–200, 207; facts 208; in- 201; normative 133, 148, 206–10, 214, 218; set-theoretical 206, 208
- simple logical autonomy thesis* 3, 19, 24, 28, 97, 113, 120, 163
- simple metaphysical autonomy thesis* 163
- Singer, D. J. 30, 35–7, 39, 89–90
- Singer's: account 35–9, 89–90; autonomy thesis 39; taxonomy 36–7
- SLAT 3, 163
- SM-irrelevance 117; and -irrelevant 117, 119–20
- SM-parts 120
- SM-states 120
- SMAT 163, 167–8, 171
- social sciences 211
- Socrates 52, 59, 167, 207
- specificity 15–6; and -desideratum 27, 72, 74
- spurious counterexamples 3–7, 46, 70, 83, 86–8, 90–1, 97–8, 102, 109–10, 114, 116, 120–2, 124, 128, 151–2, 163, 176, 180–1, 184–6, 194, 197–9, 201, 214, 217–8
- stance-independent 139
- State Gap 121, 202–3
- State RAT 183, 186–7, 202–3
- state space 52, 66, 102–5, 112, 119–21, 147, 149–50, 202–3, 205; and toy state spaces 104–5; *see also* modalized state space
- state-based account 69, 98, 110, 114, 116, 124, 147
- Sticky Situation 134–6, 139, 141–3, 146–7, 150; and Post Sticky Situation 143

- strictly normative proposition 57, 61–3; states 56, 203
- Strong Supervenience* 29
- strongly fundamental 193
- strongly normative proposition 57, 61, 66
- subject matter 14, 17, 26, 38, 43, 45–6, 110, 116–17, 120, 124, 173–4; -based 7, 14–15, 26, 43, 114, 117; -related 116, 164; -unrelated 124
- subset view 135–6
- substance-constraints 137
- substitution-test 109
- sufficiency 103
- Super-Mean 43
- supervenience 29, 37–9, 201
- Suspect Objection 141
- synthetic: account 45, 85, 105, 107, 110, 117–18, 197; and taxonomy 108
- T*-confirmer 106
- Tactful/TACTFUL 129, 132, 137–9
- taxonomic desiderata 15
- taxonomic principles 17–19, 23, 39, 173, 217
- Theory of language 145
- thick normative concepts 55, 129, 132, 210–11
- thick-thin distinction 133
- thin normative concept 128–30, 132–3
- Thin Relevance* 133
- Titelbaum, M. G. 213
- torture 36–37, 170
- toy-example 213
- transitive 28, 52, 162
- translation-fragility 31–2
- Treatise on Human Nature 1
- Trogon, K. 192
- truth-apt 13
- truth: -conditions 39, 45, 73, 91, 130, 148; -evaluation 33, 38, 9; -functional compounds 17; -preservation 109, 201; and -value 7, 30–1, 33, 35–6, 39–40, 42–3, 109, 187
- truthmaker: -based account 120, 217; -based analysis 66, 72, 194; -based approach 6; -based interpretation 83; -based semantics 183; framework 5, 8, 54, 110, 149, 194, 202, 214, 218; semantics 5, 8, 51, 59, 70, 74, 84, 181–2, 202
- truthmaking 111
- Truth-Teller 5, 75; and Truth-Teller* 53
- unconnectedness 109
- unified: account 4, 87, 135, 199; domain of normativity 7; explanatory approach 217
- unilateral account of propositions 53
- universal proposition 168
- unrestricted inference barriers 86–8
- Upwards-N-Closure 55, 66, 84
- utilitarian theory 149
- vacuous: entailments 18; and occurrences 82
- value-free science 210
- variance 30; and -based taxonomies 39, 70
- Väyrynen, P. 207
- Vranas, P. 34
- Weak Full Ground 182; and Weak Full Ground* 183
- weakly* fundamental 193
- weakly normative proposition 57–8, 61, 66, 75, 181
- What-to-do-questions 134, 136–8, 140–2, 145–6, 148, 150
- why- and whether-questions 164
- Witmer, D. G. 192
- Wodak, D. 200
- Woods, J. 20, 34, 36–7, 204
- World-Autonomy 95
- world-based: accounts 46; and taxonomies principles 38
- world-norm pair 35–39
- Xanthippe 202
- zero-ary 107