

Value Incommensurability and Practical Reason

By

David Moon

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy
(Philosophy)

at the

UNIVERSITY OF WISCONSIN-MADISON

2018

Date of final oral examination: 4/25/2018

This dissertation is approved by the following members of the Final Oral Committee:

Russ Shafer-Landau, Professor, Philosophy

Daniel M. Hausman, Herbert A. Simon Professor, Philosophy

Peter Vranas, Professor, Philosophy

Michael G. Titelbaum, Associate Professor, Philosophy

Paul Kelleher, Associate Professor, Medical History and Bioethics

Contents

Acknowledgements.....	ii
List of Acronyms and Abbreviations.....	iii
Abstract.....	iv
Chapter 1: Ontological Incommensurability.....	6
1.1 – Soft Comparativism	6
1.2 – What Grounds Deontic Statuses in OI Cases: Three Theories	12
1.2.1 – The Optimizing Norm.....	12
1.2.2 – The Maximizing Norm	16
1.2.3 – The Optimizing + Parity Norm.....	18
1.3 – Incommensurability as Indeterminacy.....	21
Chapter 2: Epistemic Incommensurability.....	26
2.1 – Epistemic Incommensurability Cases	26
2.1.1 – Recap of Chapter 1.....	26
2.1.2 – EI Cases	28
2.2 – Maximizing Actual Value.....	33
2.3 – Maximizing Expected Value and the Principle of Insufficient Reason.....	39
2.3.1 – The MEV Norm	39
2.3.2 – The Principle of Insufficient Reason	44
2.3.3 – The Multiple Partitions Objection	47
2.3.4 – The Interval Scale Objection	53
2.3.5 – The Equiprobability Objection	59
2.4 – Maximin.....	67
2.4.1 – The MAPV Norm, Maximin, Maximax, and Medianmax.....	67
2.4.2 – Why Maximin is the Operative Norm in EI Cases	73
Chapter 3: Apparent Incommensurability.....	94
3.1 – OI-EI Uncertainty	94
3.1.1 – Introduction.....	94
3.1.2 – It Is Difficult (Perhaps Impossible) To Distinguish OI and EI Cases.....	98
3.2 – OI-EI Uncertainty Affects Subjective Deontic Statuses Via the Stringency Norm.....	116
3.3 – The Stringency Norm’s Effect on Putative OI Cases... ..	129
3.3.1 – ...Under a Maximizing Theory of OI.....	129
3.3.2 – ...Under an Optimizing Theory of OI	133
3.3.3 – ...Under a Parity Theory of OI.....	136
3.4 – Conclusion	138
Bibliography	142

Acknowledgements

I would first like to thank my advisor, Russ Shafer-Landau, for all of his help, feedback, and advice. Thanks also to the rest of my committee—especially Peter Vranas and Dan Hausman—for their helpful comments and criticism.

Although I did not work with Alan Rubel on the dissertation, he did more than anyone else to help me navigate the challenges of jointly pursuing a J.D. and Ph.D. I am grateful for all of the advice and opportunities that he provided along the way.

Thanks are also due to Jeanine Grenberg. It was through her sharp and charismatic teaching that I first fell in love with philosophy. Her early modern class turned me into a philosophy major, and her Kant seminar convinced me to pursue graduate studies in philosophy.

Finally, thanks to my wife, Alaina Moon, and our children, Sage and Rowan. I could never have finished this dissertation without their support and sacrifices.

List of Acronyms and Abbreviations

EI	epistemic incommensurability
the MAPV norm	a rational norm of maximizing average possible value
Maximax	a rational norm of maximizing maximum possible value
Maximin	a rational norm of maximizing minimum possible value
the MAV norm	a rational norm of maximizing actual value
Medianmax	a rational norm of maximizing median possible value
the MEV norm	a rational norm of maximizing expected value
OI	ontological incommensurability
OI-EI uncertainty	agential uncertainty as to whether the agent faces an OI case or an EI case

Abstract

In cases of apparent value incommensurability, an agent's alternatives appear to her to be neither better than, nor worse than, nor equally as good as one another. It is plausible to suppose that some cases of apparent incommensurability will turn out to be instances of genuine, ontological incommensurability ("OI"), while others will turn out to be cases of merely epistemic incommensurability ("EI"). In EI cases, the alternatives are in fact commensurable—they are either better than, worse than, or equally as good as one another—but the agent is ignorant as to which of these value relations holds between them.

Cases of apparent incommensurability have long posed a vexing challenge for theorists of practical rationality, because we typically suppose that an alternative is rationally permissible (absent deontic constraints) if and only if it is at least as good as each other alternative—that is, if it is (one of) the best alternative(s). But in cases of apparent incommensurability, there either is no best alternative (if the agent faces an OI case), or the agent is irremediably ignorant as to which alternative is the best (if she faces an EI case). Moreover, the agent may not even know whether she faces an OI case or an EI case: she may be "OI-EI uncertain."

In this dissertation, I present my theory of what rationality requires of agents who face a choice between apparently incommensurable alternatives. Chapter 1 covers OI cases; Chapter 2 covers EI cases; and Chapter 3 covers OI-EI uncertainty. While a fair amount has been written about OI cases, very little has been written on the practical rationality of EI cases. There is even less literature—really, none at all—on the topic of whether and how OI-EI uncertainty affects rational choice. I therefore content myself in Chapter 1 with a presentation of the main existing theories of how practical rationality operates in OI cases, and focus my original efforts on Chapters 2 and 3.

Introduction

We sometimes—arguably, often—face choices in which the available alternatives appear to be incommensurable. That is, neither appears better than the other, nor do they appear to be equally good. This appearance of incommensurability typically (though not exclusively) arises in choices with qualitatively disparate alternatives. Suppose, for example, that you can afford one non-entrée splurge at a pricey Italian restaurant, and have narrowed your choice down to the carpaccio appetizer or the cannoli dessert. The two are simply so different that it seems difficult to judge which is better, or even to say with confidence that it is the case that one *is* better. Yet it also seems difficult to say anything as precise as that they are equally good.

Cases of apparent incommensurability are not limited to such trivial choices. Nor are they limited to any particular domain of practical reason. In legal reasoning, a judge might find herself having to choose between two mutually inconsistent rulings that appear incommensurable to her with respect to the relevant legal value(s). This might happen, for example, in cases that involve an especially thorny issue of constitutional interpretation. Perhaps one possible ruling is supported by strong textual reasons and certain parts of the historical record, but seems inconsistent with an important line of cases. The other possible ruling fits well with these cases and is supported by other parts of the historical record, but seems inconsistent with a strict textualist reading of the Constitution. Or in moral reasoning, an agent might find herself facing two possible actions that appear incommensurable with respect to the relevant moral value(s). And in epistemic reasoning, an agent might confront a set of mutually exclusive propositions that appear incommensurable with respect to the relevant epistemic value(s).

All of these examples involve individual agents, but cases of apparent incommensurability can arise in collective agency as well. We may confront them when acting in

our capacities as policy makers or voters, and in other cases where we must choose as part of or on behalf of a group.

Sometimes further reflection or learning more information about one's alternatives resolves an ostensible incommensurability. But often the impression of incommensurability persists despite thorough reflection and a good grasp of all relevant empirical facts. The examples above, and others like them, remain compelling even if we stipulate that you are well-acquainted with carpaccio and cannoli, that the judge is a renowned constitutional scholar who has thoroughly examined the case's factual record, and so on.

The phenomenon of apparent, persistent incommensurability poses a problem for practical reasoning. The problem arises because we typically take the rational justification of our choices to depend on some comparative, quantitative fact or belief about the alternatives' value (e.g., that the chosen alternative is at least as good as each other alternative), perhaps in conjunction with facts about all relevant deontic constraints. But in cases of apparent, persistent incommensurability, we either cannot ascertain the comparative, quantitative fact, or there is no such fact.

In this dissertation, I offer my answer to the following question: *when forced to choose between apparently incommensurable alternatives, what should we do, and why?* I interpret this as a question about what metaphysically grounds alternatives' deontic statuses (permissibility, impermissibility, obligatoriness, etc.). If we determine which facts, beliefs, and rational norms ground alternatives' deontic statuses, we can then ask, when confronted with a case of apparent incommensurability, what deontic statuses those norms, facts and beliefs render.

It is plausible to suppose that some cases of apparent incommensurability will turn out to be instances of genuine, ontological incommensurability ("OI"), while others will turn out to be

cases of merely epistemic incommensurability (“EI”). In OI cases, the agent’s alternatives are neither better than, nor worse than, nor equally as good as one another. In EI cases, the alternatives are in fact commensurable—they are either better than, worse than, or equally as good as one another—but the agent is irremediably ignorant as to which of these value relations holds between them. OI and EI are mutually exclusive and exhaustive. Hence, all cases of apparent incommensurability are instances of either OI or EI.

Chapter 1 covers OI cases, Chapter 2 covers EI cases, and Chapter 3 covers what I call “OI-EI uncertainty.” In cases of OI-EI uncertainty, the agent is uncertain whether she faces an OI case or an EI case. Together, these three chapters comprise a complete account of what grounds alternatives’ deontic statuses in cases of apparent incommensurability. They cover each half of a mutually exclusive and exhaustive subdivision of cases of apparent incommensurability (into OI and EI cases), and also cover how practical reason operates when there is uncertainty as to which half of the subdivision applies to the choice at hand. While a fair amount has been written about OI cases, the literature on EI cases and OI-EI uncertainty is scant to nonexistent. I therefore content myself with a presentation of existing theories of OI in Chapter 1, and focus my efforts on making new, original contributions in Chapters 2 and 3.

I begin Chapter 1 by adopting a barebones theory of practical rationality that I call “soft comparativism.” According to soft comparativism, alternatives’ deontic statuses are grounded in (1) a comparative, quantitative fact or belief about the alternatives’ value; (2) a rational norm that takes the comparative, quantitative fact or belief’s propositional content as an input; and (3) any relevant deontic constraints. I then present the three main theories regarding which rational norm grounds alternatives’ deontic statuses in OI cases. The first candidate rational norm is an *optimizing* norm, which holds that an alternative is rationally permissible (absent deontic

constraints) if and only if it is *at least as good as* each other alternative. The second is a *maximizing* norm, according to which an alternative is rationally permissible if and only if it is *not worse than* each other alternative. According to the third candidate norm, which I call the *optimizing + parity* norm, an alternative is rationally permissible if and only if it is either at least as good as *or on a par with* each other alternative.

In Chapter 2, I develop a theory of what grounds alternatives' deontic statuses in EI cases. EI cases, unlike OI cases, essentially involve agential ignorance or uncertainty. One of the three standard value relations (better than, worse than, or equally good) holds between the agent's alternatives, but she is ignorant as to which. I distinguish between objective and subjective deontic statuses, and argue that a norm of *maximizing actual value* plausibly grounds the former, but not the latter. I then argue that because EI cases are better understood as decisions under ignorance rather than decisions under risk, a norm of *maximizing expected value* cannot ground subjective deontic statuses, either. I consider, but ultimately reject, the possibility that this norm might ground subjective deontic statuses in EI cases with the help of the principle of insufficient reason (sometimes called the principle of indifference). Finally, I argue that subjective deontic statuses in EI cases are grounded in a "maximin" norm of *maximizing minimum possible value*.

In Chapter 3, I develop a theory of how OI-EI uncertainty—that is, uncertainty on the part of the agent as to whether she faces an OI case or an EI case—affects subjective deontic statuses in cases of apparent incommensurability. After briefly explaining why OI-EI uncertainty does not affect *objective* deontic statuses, I begin by arguing that OI-EI uncertainty is present in most cases of apparent incommensurability. I then introduce the notion of stringency and argue that OI-EI uncertainty affects subjective deontic statuses via the stringency norm. The stringency

norm is a second-order rational norm, analogous to the first-order rational norm of dominance, that adjudicates between norm-belief pairs. In cases of apparent incommensurability, the stringency norm adjudicates between (1) the norm and belief that would jointly ground subjective deontic statuses if the choice were known to be an OI case, and (2) the norm and belief that would jointly ground subjective deontic statuses if the choice were known to be an EI case. I finish by examining how the stringency norm affects subjective deontic statuses in cases of apparent incommensurability under each of the three theories of OI from Chapter 1. Under all three theories, the presence of OI-EI uncertainty and the operation of the stringency norm have the potential to affect subjective deontic statuses in putative OI cases.

Chapter 1: Ontological Incommensurability

1.1 – Soft Comparativism

In this dissertation, I present my solution to a particular philosophical puzzle. The puzzle, roughly stated, is to satisfactorily answer the following question: what does rationality require of agents facing a choice between alternatives that appear to her (and to us, as third-party observers) to be incommensurable? Put differently: when forced to choose between apparently incommensurable alternatives, what should we do, and why? To solve the puzzle—and indeed, to even state the puzzle more precisely than this—I will need to make some basic assumptions about practical rationality. In this section, I spell out those assumptions. My aim is to adopt a barebones theory of practical rationality that has as few substantive commitments as possible.

In this dissertation, I discuss the rationality of *choices* made by an *agent* between mutually exclusive *alternatives*. “Agent” should be liberally construed to include both individuals and groups of individuals, making choices either on their own behalf or in some other capacity (e.g., as a voter, policymaker, judge, or proxy). Although there are important differences between various types of agents, none of my claims will hinge on those differences. I take a similarly catholic approach in defining “alternative.” Some philosophers prefer to describe the agent’s available alternatives as objects to be chosen (e.g., carpaccio) or acts to be performed (e.g., ordering carpaccio). Others insist that alternatives are best understood as states of affairs to be brought about (e.g., bringing about a state of affairs that includes my eating carpaccio). I will remain neutral in the debate on how to best characterize alternatives.

Alternatives have *deontic statuses*: they may be rationally permissible, impermissible, obligatory, etc. Those deontic statuses have a *metaphysical ground*. There is an ever-growing

literature on metaphysical grounding.¹ Here, too, I will try to remain as neutral as possible. We might say that there is something *in virtue of which* an alternative has the deontic status that it does, or that there is something which *makes it the case that* an alternative has a particular deontic status, or that there is something which *non-causally explains* why an alternative has some deontic status rather than another, or that there is something that serves as the *truthmaker* for the proposition that an alternative is (for example) rationally permissible, and so forth. Flesh out the particulars how you like. I will have a lot to say about what metaphysically grounds alternatives' deontic statuses, but my claims will not be inextricably linked to any particular theory of metaphysical grounding.

Here is an influential thesis about what metaphysically grounds alternatives' deontic statuses:

comparativism: An alternative's deontic status is metaphysically grounded in: (1) a comparative, quantitative fact about the alternatives' value (e.g., the fact that *A* is at least as good as *B*); and (2) a rational norm that takes the comparative, quantitative fact's propositional content as an input (e.g., a norm according to which an alternative is rationally permissible if and only if it is at least as good as each other alternative).

“Quantitative” here may be understood liberally. It need not imply anything as strong as representability on a cardinal scale or an ability to assign precise, numeric values.

Much of the contemporary debate on value incommensurability assumes some variant of comparativism. This is partly due to the fact that Ruth Chang, who is one of the most prolific and influential writers on this topic, holds and defends a comparativist view of practical reason. (It was Chang who coined the term “comparativism” (see Chang 1997b, 9).) But comparativism is not original to Chang—she herself likens it to an “old and delicious wine” (2016a, 216).

¹ For an overview, see Bliss & Trogon 2016.

According to comparativism, it is facts about alternatives' comparative value, together with rational norms that take the propositional content of those facts as inputs, that rationally justify our actions. Comparativism has deep intuitive appeal. When I make a choice—be it a moral choice, a prudential choice, or a choice in some other domain of practical reason—we typically suppose that I ought to choose the best alternative, and that it would be irrational of me to choose an alternative that is worse than another available alternative. (Equivalently: the worse alternative is rationally impermissible.) Moreover, we take these value-comparative facts—that *A* is the best alternative, that *B* is the worst alternative, and so on—to metaphysically ground the alternatives' deontic statuses. I ought to choose the best alternative (or abstain from choosing a worse alternative) in virtue of the value-comparative fact that it is the best alternative.

Despite its intuitive appeal, many of us may nevertheless find Chang's version of comparativism unpalatable. Because it grounds deontic statuses solely in facts, and never in beliefs, it applies only to objective obligations (and other objective deontic statuses), and not to subjective obligations (and other subjective deontic statuses). But subjective deontic statuses—which reflect the rationality of our choices relative to our (potentially incomplete or mistaken) beliefs about our circumstances, rather than relative to our actual circumstances—are also an important component of practical rationality.² Some theorists go so far as to argue that only subjective deontic statuses are genuine, and that there is no need to posit objective deontic statuses at all (see Smith 2010b, 70). But even if we retain objective deontic statuses, so long as we think (as I do) that subjective deontic statuses have at least some role to play in a theory of practical rationality, the version of comparativism formulated above must be viewed as an incomplete statement of the metaphysical grounds of deontic statuses.

² For a more thorough description of the difference between subjective and objective deontic statuses, see Chapter 2, section 2.2.

Another shortcoming of comparativism, as formulated above, is its failure to account for *deontic constraints*. Deontic constraints make it the case that an alternative is rationally impermissible even if the relevant value-comparative fact/belief and rational norm would otherwise render it rationally permissible. In the moral domain, for example, an alternative that is at least as good as every other alternative may nevertheless be impermissible because it violates what Robert Nozick (1974) calls a “side constraint.” Or it may be impermissible on the basis of what Joseph Raz (1999, 35-48) calls an “exclusionary reason.”³

In the interest of making our barebones theory of practical rationality as inclusive and uncontroversial as possible, we can incorporate these elements into a softer version of comparativism:

soft comparativism: An alternative’s deontic status is metaphysically grounded in: (1) a comparative, quantitative fact *or belief* about the alternatives’ value; a (2) a rational norm that takes the comparative, quantitative fact *or belief*’s propositional content as an input; *and* (3) *any relevant deontic constraints*.

Soft comparativism is something that I think strikes most of us as plausible. It plausibly applies to both trivial decisions (e.g., which breakfast cereal to buy) and terribly weighty ones (e.g., whether to go to war). With a sufficiently inclusive interpretation of “value,” soft comparativism also applies across different domains of practical reason: prudential, epistemic, legal, moral, etc. Even proponents of “hard” comparativism can accept soft comparativism, and maintain either that the portions in italics have nothing in their extension (e.g., by maintaining that there are never any relevant deontic constraints), or that those things picked out by the portions in italics

³ Jeremy Waldron (2003, 196) provides a helpful example of an exclusionary reason: “Consider...the reasons associated with promise-keeping. If I have already promised to meet with a student to discuss his paper at 12.30 p.m., then I may not accept an invitation to lunch with a colleague at that time.” Even if I later “find out that this is going to be a *really delicious lunch*[,] ... [t]he existence of the promise provides a reason for not acting on considerations like the quality of the lunch.”

are reducible to the non-italicized portions (e.g., by maintaining that it is always possible to re-describe any relevant deontic constraints as value-comparative facts⁴). Throughout the dissertation, I will assume that soft comparativism is true.

Absent deontic constraints, soft comparativism grounds alternatives' deontic statuses in facts or beliefs about their *value*. I will construe "value" broadly to mean something like "consideration relevant to the choice." More precisely, I define a value as *anything in terms of which a value relation* (better than, worse than, or equally good) *can hold*. Under this definition, examples of "values" include not only considerations such as generosity and parsimony, but may even include metrics from the natural sciences. For example, I might go to Home Depot with the purpose of buying the tallest ladder I can find. Once there, I would compare my alternatives in terms of length: a taller ladder is better with respect to length than a shorter ladder.

I will assume that for any choice, there is a value that is relevant to the choice. It is facts or beliefs about this value—the relevant value—that ground the alternatives' deontic statuses. In most choices, the relevant value is itself a composite of several values. The value relevant to my choice of which house to buy, for example, is a composite value that includes affordability, attractiveness, neighborhood walkability, and many other constituent values. The relevant value is an overarching value that includes these constituent values.

I will say that alternatives are *incommensurable* with respect to the relevant value if and only if they are neither better than, nor worse than, nor equally as good as one another with

⁴ Ruth Chang (2016a, § 3.1) is a hard comparativist who holds this view about deontic constraints.

respect to the relevant value.⁵ Equivalently: alternatives are incommensurable if and only if none of the three standard value relations (better than, worse than, equally good) holds between them.

This is a dissertation about cases of *apparent value incommensurability*. In cases of apparent incommensurability, the alternatives appear to the agent to be incommensurable, despite thorough reflection and a good grasp of all relevant empirical facts. Cases of apparent incommensurability subdivide into two mutually exclusive and exhaustive categories. Some cases of apparent incommensurability are cases of genuine, ontological incommensurability (“OI”). In OI cases, the alternatives are incommensurable. Other cases of apparent incommensurability are cases of merely epistemic incommensurability (“EI”). In EI cases, the alternatives are in fact commensurable. They are either better than, worse than, or equally as good as one another, but the agent is ignorant as to which value relation holds between them.

Cases of apparent incommensurability may be either *complete* or *partial*. In cases of complete apparent incommensurability, every alternative appears incommensurable with every other alternative. In cases of partial apparent incommensurability, on the other hand, there is at least one pair of alternatives that appear commensurable with respect to the relevant value (and at least one pair of alternatives that appear incommensurable).

Having incorporated the notion of deontic constraints into soft comparativism, I will now, sadly, set them aside. There is much to be said on the topic of how and when deontic constraints ground deontic statuses in cases of apparent incommensurability.⁶ But in this dissertation, I will focus instead on the question of what metaphysically grounds alternatives’ deontic statuses in cases of apparent incommensurability in which there are no relevant deontic constraints.

⁵ When I make an assertion of the form “*A* is better than *B*,” I should always be read as saying that “*A* is better than *B* with respect to the relevant value.” For brevity, I will often omit “with respect to the relevant value.”

⁶ I have written a bit about this topic elsewhere (Moon, unpublished b).

Answering the question of what grounds alternatives' deontic statuses will also give us the means to answer the further question of what those deontic statuses are.

I will begin in this chapter by presenting the three leading theories of what grounds alternatives' deontic statuses in OI cases. My goal in the rest of Chapter 1 is to present a quick summary of ideas and arguments from the existing literature on the practical rationality of OI cases. I will remain neutral on the question of which of the three theories of OI is correct. One reason for my neutrality is that OI cases have already been relatively well discussed, and I want to focus my efforts on areas that have not received the same attention. One of these neglected topics, which will be the subject of Chapter 2, is the question of what grounds alternatives' deontic statuses in EI cases. The other neglected topic, which will be the subject of Chapter 3, is the question of how OI-EI uncertainty—that is, uncertainty on the part of the agent as to whether she faces an OI case or an EI case—affects alternatives' deontic statuses.

My choice of topic in Chapter 3 is the other reason for my neutrality, in Chapter 1, as to which theory of OI is correct. Once I present my theory (in Chapter 3) of how OI-EI uncertainty affects alternatives' deontic statuses, I will present what my theory's implications are under each of the three theories of OI from Chapter 1. By remaining neutral on the question of which theory of OI is correct, I ensure that the usefulness of my work in Chapter 3 (on OI-EI uncertainty) is not tied to any particular theory of OI.

1.2 – What Grounds Deontic Statuses in OI Cases: Three Theories

1.2.1 – The Optimizing Norm

Absent deontic constraints, alternatives' deontic statuses are jointly grounded in (1) a comparative, quantitative fact or belief about the alternatives' value, and (2) a rational norm that takes the fact or belief's propositional content as an input. In OI cases, the relevant fact or belief

is that the alternatives are ontologically incommensurable: they are neither better than, nor worse than, nor equally as good as one another. When discussing OI cases in this chapter, I will assume that the agent knows that she faces an OI case, rather than an EI case. (I will take up the question of OI-EI uncertainty later, in Chapter 3.) If the agent knows that she faces an OI case, then she both believes that her alternatives are ontologically incommensurable, and it is a fact that her alternatives are ontologically incommensurable.

There are three main theories as to which rational norm grounds alternatives' deontic statuses in OI cases. The first theory is that alternatives' deontic statuses in OI cases are grounded in an *optimizing* norm. According to the optimizing norm, an alternative is rationally permissible (absent deontic constraints) if and only if it is *at least as good as* each other alternative. Equivalently: an alternative is rationally permissible (absent deontic constraints) if and only if it is the *best* alternative.⁷

Conventional wisdom holds that the optimizing norm typically grounds alternatives' deontic statuses.⁸ Indeed, absent deontic constraints, how could it be rational to knowingly choose anything other than the best alternative? If I am buying a house, for example, it is surely rationally impermissible for me to choose one that I know is worse with respect to the relevant value—say, more expensive, uglier, and in a crummier neighborhood—than another that is for sale. I ought choose the best of my available alternatives, on pain of irrationality.

⁷ “Best” here does not mean “uniquely best.” Two alternatives, each of which is better than any alternative save each other, may be equally good.

⁸ Ruth Chang suggests that the optimizing norm “is explicit in instrumentalist approaches to rational choice and decision theory, and implicit in most forms of consequentialism, some versions of virtue theory, and, arguably, certain forms of deontology” (1998, 1577-1578).

If an agent's set of alternatives is fully commensurable,⁹ there will always be at least one alternative that is the best, with respect to the relevant value. But this is not true of the alternatives in OI cases. If an agent faces a set of ontologically incommensurable alternatives, then no alternative is at least as good as each other alternative. And if an alternative's being at least as good as each other alternative is a necessary and sufficient condition for its rational permissibility (as the optimizing norm holds), it follows that no alternative is rationally permissible. Hence, if the optimizing norm is operative in OI cases, practical reason breaks down. An agent faces a set of alternatives, none of which is rationally permissible. Yet she must choose one of them.

Normally, we would take the claim, entailed by the optimizing norm, that no alternative in an OI case is rationally permissible, to mean that every alternative is rationally *impermissible*. This follows from the normally innocuous-seeming assumption that permissibility and impermissibility are mutually exclusive and exhaustive. But retaining that assumption in the context of OI cases leads to the violation of another normally innocuous-seeming assumption that it is never the case that every alternative is rationally impermissible. It can be argued that both of these assumptions are conceptual truths that follow from the concepts of rational permissibility and rational impermissibility. Yet taken together with the optimizing norm in the context of an OI case, they yield a contradiction: it is never the case that every alternative is rationally impermissible, and in OI cases every alternative is rationally impermissible.

Proponents of the optimizing norm are best characterized as avoiding this contradiction by partially rejecting the first of the two innocuous-seeming assumptions. Specifically, they reject that permissibility and impermissibility are *exhaustive*. (They need not, and do not, reject

⁹ By "fully commensurable," I mean that each alternative is commensurable with every other alternative.

the second half of the assumption: viz., that permissibility and impermissibility are mutually exclusive.) Permissibility and impermissibility are not exhaustive, on their view, because there is a third logical possibility that obtains in OI cases: there might simply be no deontic status at all.

If permissibility and impermissibility are not exhaustive, then the claim that no alternative in an OI case is rationally permissible no longer entails that every alternative is *impermissible*. It might instead be the case that each alternative has *no* deontic status. This latter possibility—that alternatives have no deontic status at all, rather than having the deontic status of rational impermissibility—is the one that obtains in OI cases, according to proponents of the optimizing norm. The optimizing norm fails to render any deontic statuses at all in OI cases, and practical reason is silent. Practical reason simply has no guidance to offer to the agent, and she is forced to choose by “plumping”—that is, arbitrarily selecting an alternative with no guidance from practical reason.¹⁰

The optimizing norm states necessary and sufficient conditions for rational permissibility: an alternative is permissible (absent deontic constraints) if and only if it is at least as good as each other alternative. But there are also (the optimizing norm’s proponents might say) necessary and sufficient conditions for rational *impermissibility*: an alternative is *impermissible* (absent deontic constraints) if and only if it is *worse than* each other alternative. In OI cases, alternatives are neither at least as good as nor worse than one another. Just as they fail to meet the necessary and sufficient conditions for rational permissibility, so too they fail to meet the necessary and sufficient conditions for rational impermissibility. Reasons “run out,” and “[practical] reason fail[s] to adjudicate between” the agent’s alternatives (Raz 1997, 125). Rather than making an *irrational* choice (as she would if every alternative were impermissible), the agent makes an

¹⁰ For a helpful discussion of how “plumping” differs from the “picking” that agents do when they select one of several rationally permissible alternatives, see Chang 2017, § 3.

arational choice (because practical reason fails to give her alternatives any deontic statuses at all).¹¹

If the optimizing norm is operative in OI cases, then OI cases represent a failure of practical rationality. When alternatives are ontologically incommensurable, they have no deontic statuses, and the agent is left to plump between them.

1.2.2 – *The Maximizing Norm*

Some philosophers (e.g., Sen 1997 & 2000; Hsieh 2007) hold that in OI cases, alternatives' deontic statuses are grounded not in an optimizing norm, but instead in a *maximizing* norm. Under the maximizing norm, an alternative is rationally permissible (absent deontic constraints) if and only if it is *not worse than* each other alternative.

If a set of alternatives is fully commensurable, the set of alternatives that maximizes is identical to the set of alternatives that optimizes. But when the alternatives are incommensurable—i.e., in OI cases—the maximizing and optimizing norms come apart. If an agent's alternatives are incommensurable, then they are not worse than one another.¹² Under the maximizing norm, every member of a set of incommensurable alternatives is therefore rationally permissible.

My sense, in speaking with other philosophers on this topic, is that most of us—at least, those who don't suspend judgment on the matter—have intuitions that favor either optimizing or maximizing. Those who think that the maximizing norm is insufficient to rationally justify choice emphasize that the optimizing norm grounds deontic statuses in the fact that some value relation (either better than or equally good) holds between the alternatives. The maximizing

¹¹ I borrow this handy terminological distinction from Chang 2017, § 3.

¹² Incommensurability consists in alternatives' being neither better than, *nor worse than*, nor equally as good as one another.

norm, in contrast, purports to ground deontic statuses not in a value relation that holds between the alternatives, but in the absence of one. Ruth Chang, for example, finds the premise that only a “positive value relation” may ground deontic statuses to be sufficiently self-evident that she states without argument that it is an underlying assumption of her version of comparativism (see, e.g., 1997b, 4; 2002b, 663; 2012, 112; 2016a, 217). Those who favor the maximizing norm often take it to be equally self-evident that this premise is false. Why, they say, should it matter to the maximizing norm’s plausibility that it states a “negative” rather than “positive” relation between the alternatives?

Nien-hê Hsieh (2007) argues that the maximizing norm is more intuitively plausible than the optimizing norm, and is thus presumptively correct. The maximizing norm’s plausibility, says Hsieh, lies in its “key intuition” that “the choice of a worse alternative is unjustified” (2007; 72). Because the maximizing norm is presumptively correct, Hsieh argues, the burden of proof is on those who would reject it.

But proponents of the optimizing norm can reply that Hsieh’s “key intuition” is simply a restatement of the impermissibility condition discussed in the last section: an alternative is rationally impermissible (absent deontic constraints) if and only if it is not worse than each other alternative. This impermissibility condition, optimizers can argue, is consistent with the optimizing norm’s claim that alternatives are rationally permissible if and only if they are at least as good as each other alternative. It is, moreover, an integral part of the more plausible of the two possible interpretations of the optimizing norm’s verdict in OI cases. Rather than supporting the maximizing norm over the optimizing norm, they can argue, it simply restates part of the optimizing theory of OI.

The optimizing norm entails that in OI cases, no alternative is rationally permissible. The maximizing norm avoids this unfortunate result by grounding deontic statuses in something other than the fact (or belief) that a value relation holds between alternatives—specifically, in the fact that a particular value relation, “worse than,” *fails to hold* between alternatives. Some philosophers find the maximizing norm to be intuitively plausible. Other philosophers believe that deontic statuses must be grounded in a “positive value relation.” If the latter group hopes to avoid the optimizing norm’s unfortunate result, they must instead posit a new, nonstandard value relation (i.e., a value relation that is distinct from “better than,” “worse than,” or “equally good”) that holds between alternatives in OI cases. This is the strategy employed by proponents of the third theory of OI, to which I now turn.

1.2.3 – *The Optimizing + Parity Norm*

According to conventional wisdom, there are three value relations that may hold between alternatives: better than, worse than, or equally good. Two alternatives are incommensurable just in case none of these three value relations holds between them. In the past twenty years or so, several philosophers¹³ have argued that there is a fourth value relation called “parity,” or “on a par,” that sometimes holds between incommensurable alternatives. According to proponents of parity, the operative rational norm in OI cases is neither the optimizing norm nor the maximizing norm. They maintain that OI cases (and non-OI cases) are instead governed by what I’ll call the “optimizing + parity” norm. Under the optimizing + parity norm, an alternative is rationally permissible if and only if it is *either* at least as good as *or on a par with* each other alternative. Thus, two (or more) on-a-par alternatives have the same deontic statuses under the optimizing +

¹³ Parity’s main proponents are Ruth Chang (2002a, Ch. 5-6; 2002b; 2005; 2016b), Wlodek Rabinowicz (2008; 2009a; 2009b; 2012), Joshua Gert (2015).

parity norm that incommensurable alternatives have under the maximizing norm: either is rationally permissible.

Like alternatives that are equally good, alternatives that are on a par are neither better than nor worse than one another. And like equality, parity is symmetric: if *A* is on a par with *B*, then *B* is on a par with *A* (Chang 2016b, 195). But unlike equality, parity is neither reflexive (“*A* is never on a par with itself: the two are equally good”) nor transitive (“if *A* is on a par with *B*, and *B* is on a par with *C*, then it does not follow that *A* is on a par with *C*”) (Chang 2016b, 195). Alternatives that are on a par therefore cannot be ordinally ranked with one another with respect to the relevant value, because ordinality presupposes transitivity.¹⁴

Alternatives that are on a par are not equal. Instead, they are “in the same neighborhood” of value (Chang 2016b, 193). But parity should not be confused with EI cases, in which one of the three standard value relations holds between the agent’s alternatives, and she is ignorant as to which relation holds. The judgment that two alternatives are on a par does not result from partial ignorance about alternatives’ relative value. Parity is instead like James Griffin’s similar notion of “rough equality,” in which “the roughness is not in our understanding but ineradicably in the values themselves” (1986, 81).¹⁵ Parity is, its supporters say, an ontologically distinct value relation.

¹⁴ More specifically, transitivity is one of four necessary and sufficient conditions for constructing an ordinal scale. The other three conditions are completeness, asymmetry, and negative transitivity (Peterson 2009, Ch. 5).

¹⁵ Griffin’s rough equality resembles Chang’s parity relation in that (1) it is intransitive, and (2) alternatives that are roughly equal are neither better than nor worse than one another. But rough equality is not identical to parity. Griffin states that rough equality is an instance of “vague [ordinal] ordering” (1986, 81), while Chang denies a vagueness account of incommensurability (2002b). Parity is also similar to, but not identical with, Derek Parfit’s notion of “imprecise comparability.” For a discussion of how parity differs from imprecise comparability, see Chang 2016b.

Parity remains a controversial value relation. It is a nontransitive, irreflexive relation that fails to cohere neatly with the three standard value relations. The assumption that the three standard value relations are logically exhaustive is a long-standing tenet of decision theorists (see, e.g., von Neumann & Morgenstern 1944; Luce & Raiffa 1957; Resnik 1987; Peterson 2009), and “has all the appeal of an analytic truth” (Wasserman 2004, 392). Proponents of the maximizing norm are likely to view parity as an unnecessarily complicated way of achieving the same practical result as the maximizing norm: the rational permissibility of every member of a set of incommensurable alternatives. Nien-hê Hsieh (2005) proposes a deflationary account, according to which on a par alternatives are equally good.¹⁶ Other authors criticize either the “Small Improvement Argument” (Anderson 2015; Klocksien 2010; Espinoza 2007; Gustaffson & Espinoza 2010; Gustaffson 2013; Wasserman 2004) or the “Chaining Argument” (Andersson 2016; Boot 2009; Elson 2014), whose conclusions are both premises in Ruth Chang’s argument for the existence of parity in her (2002b) paper, “The Possibility of Parity.”¹⁷

But parity has its defenders, too. Wlodek Rabinowicz (2008; 2009a; 2009b; 2012) has incorporated it into his own preference-based theory of practical rationality. And Rabinowicz and Chang together convinced Joshua Gert, who was initially skeptical of parity (see Gert 2004), to accept and defend parity as a genuine value relation (Gert 2015, 249).

These are the three leading theories of which rational norm is operative in OI cases. If the optimizing norm is operative in OI cases, then the agent’s alternatives have no deontic statuses.

¹⁶ I argue elsewhere that Hsieh’s argument does not succeed (Moon, unpublished a).

¹⁷ The “Small Improvement Argument” purports to show that agents sometimes face alternatives that are incommensurable. The “Chaining Argument” purports to show that incommensurable alternatives are sometimes nevertheless *comparable*. Given Chang’s assumption that alternatives are comparable if and only if a positive value relation holds between them, she concludes that there must therefore be a fourth, nonstandard value relation that holds between these incommensurable-but-comparable alternatives (2002b).

If either the maximizing norm or the optimizing + parity norm is operative in OI cases, then every member of a set of incommensurable alternatives is rationally permissible. It is tempting to conclude that while OI cases present a conundrum for proponents of the optimizing norm, proponents of the other two norms already have a satisfactory resolution to the question of what rationality requires of an agent who faces an OI case. But this overlooks the fact that an agent facing ontologically incommensurable alternatives may be uncertain as to whether she faces an OI case or an EI case: she may be OI-EI uncertain. I will take up the question of how OI-EI uncertainty affects alternatives' deontic statuses in Chapter 3. Before I can do so, I need to provide an account of what grounds alternatives' deontic statuses in EI cases. This will be my task in Chapter 2.

1.3 – Incommensurability as Indeterminacy

I finish Chapter 1 by noting how the three rational norms from section 1.2 would operate under an “indeterminist” interpretation of incommensurability. Two alternatives are incommensurable if they are neither better than, nor worse than, nor equally good. Up to this point, I have taken this to mean that it is *false* (rather than true) that any of these three value relations obtains. John Broome (1997) and Luke Elson (2017) have argued that it is instead *neither true nor false* that any of the three standard value relations obtains between incommensurable alternatives. Broome and Elson (1) argue that cases of apparent incommensurability are instances in which the three standard value relations are vague; and (2) accept an indeterminacy (as opposed to epistemic) theory of vagueness. That is, they believe that if it is vague that *A* is *F*er than *B*, where *F* is some comparative predicate (e.g., better than), then it is neither true nor false that *A* is *F*er than *B*.

The operation of the optimizing norm is not affected by reinterpreting incommensurability as a case of semantic indeterminacy. The optimizing norm states that an alternative is rationally permissible (absent deontic constraints) if and only if it is *at least as good as* each other alternative. In other words, an alternative is rationally permissible (absent deontic constraints) if and only if it is *true* that that alternative is at least as good as each other alternative. Under an indeterminist interpretation of OI cases, it is neither true nor false that incommensurable alternatives are at least as good as one another. Thus, the result is the same as when we assumed it to be false that incommensurable alternatives were at least as good as one another: neither of two (or more) incommensurable alternatives is rationally permissible. Just as before, the optimizing norm fails to yield any deontic statuses, and the agent is left to mere plumping.

There are two ways that the maximizing norm might operate under an indeterminist account of incommensurability, depending on how we interpret “not worse than.” (The maximizing norm, remember, holds that an alternative is rationally permissible if and only if it is *not worse than* each other alternative.) First, we might take “*A* is not worse than *B*” to mean that it simply *fails to be true that A is worse than B*. On this reading, the maximizing norm operates in the same way that it did under our initial, truth-binary interpretation of incommensurability. If an agent faces two incommensurable alternatives, it is neither true nor false (under an indeterminist account of OI) that they are worse than one another. Hence, it fails to be true that either is worse than the other, and they are both (under this first reading of “not worse than”) rationally permissible. This is consistent with Elson’s account of deontic statuses in OI cases under his indeterminist account of incommensurability. He argues that “[w]hen two options are incommensurate ... one may permissibly choose either” (2017, 343).

Second, we might instead interpret “*A* is not worse than *B*” to mean that it is *false* that *A* is worse than *B*. On this second reading, assuming an indeterminist account of incommensurability, neither of two incommensurable alternatives maximizes. It is not false, but is instead neither true nor false, that either is worse than the other. On this interpretation of “not worse than,” the maximizing norm operates differently than it did when we assumed a truth-binary account of incommensurability. Rather than rendering both alternatives permissible, it yields the same results as the optimizing norm: it fails to render any deontic statuses at all, and the agent is left to plump.

Parity, the rationale underlying the third rational norm (optimizing + parity), is inconsistent with an indeterminist account of incommensurability. Proponents of parity propose that incommensurable alternatives are on a par. But parity is inconsistent with an indeterminist account of incommensurability, because parity is mutually exclusive with the three standard value relations. If it is true that two alternatives are on a par, then it is *false*—*not* neither true nor false—that they are either better than, worse than, or equally as good as one another. Thus, if alternatives in OI cases are on a par, then the indeterminist interpretation of OI cases is incorrect.

Finally, John Broome (1997), in presenting his indeterminist account of OI, offers an observation that might serve as the basis for a fourth candidate rational norm that operates in OI cases. To describe Broome’s account, it is helpful to imagine a chain of qualitatively similar items that are ordinally ranked with respect to the relevant value—for example, a set of poems ranked with respect to literary merit. (I will call these “type-*A* items,” and denote token type-*A* items with the variable “*A*.”) Now suppose that we are comparing a qualitatively dissimilar item (a novel, say) to each type-*A* item in the chain (see Broome 1997, 68-70). (Call this a “type-*B* item,” denoted with the variable “*B*.”) On Broome’s account, type-*A* items that are

incommensurable with *B* occupy a “zone of indeterminacy.” For any type-*A* item in the zone of indeterminacy, it is neither true nor false that any of the three standard value relations holds between it and *B*. Hence, writes Broome, “[w]hen it is indeterminate which of two things is better, their goodness is in a sense incomparable” (1997, 89).

“But,” he goes on, “in other senses [alternatives’ comparative goodness] is not incomparable” (1997, 89). Broome’s particular indeterminist theory of vagueness employs degrees of truth. A vague instance of a comparative, though not determinately true or false, is true (or false) to some degree. As we move up through the zone of indeterminacy (i.e., as we move in ascending order through the ranking of type-*A* bearers), “This poem is better than that novel” becomes *progressively truer* and “That novel is better than this poem” becomes *progressively less true*.¹⁸

This suggests that if *A* is near the top of the zone of indeterminacy, “*A* is better than *B*” is *truer than* “*B* is better than *A*”—even though, strictly speaking, “*A* is better than *B*” is not actually true. It is not hard to imagine how one might develop and defend a view that this sort of proposition—namely, a proposition of the form “*p* is truer than *q*,” where *p* and *q* are both propositions stating that a value relation holds between alternatives—might serve as the basis for a rational norm that operates in OI cases.

This move (i.e., defending a “*p* is truer than *q*” rational norm as operative in OI cases) is unavailable to Broome. A key premise in his argument for an indeterminist account of OI is the “collapsing principle,” which states that “[f]or any *x* and *y*, if it is more true that *x* is *F*er than *y* than that *y* is *F*er than *x*, then *x* is *F*er than *y*” (1997, 77). (*F* here represents any two-place

¹⁸ Broome refers to this sort of indeterminacy as “type (c) soft indeterminacy” (1997, 79-81). For Broome’s arguments that this is the correct account of indeterminacy in vague applications of comparatives, see 1997, 81-88.

comparative.) The collapsing principle entails that for all type-*A* items in the zone of indeterminacy, it cannot be the case that “*A* is better than *B*” is truer than “*B* is better than *A*” (or vice versa). If that were true, then from the collapsing principle it would follow that “*A* is better than *B*.” And if *A* is better than *B*, then *A* cannot be in the zone of indeterminacy. In short, (1) the collapsing principle, (2) Broome’s account of vague value relations, and (3) the proposition “ ‘*A* is better than *B*’ is truer than ‘*B* is better than *A*’ ” jointly entail a contradiction. Because Broome is committed to (1) and (2), he must reject (3).

But the collapsing principle has been widely criticized (see, e.g., Carlson 2004; Elson 2014a; Espinoza 2008; Gustafsson 2013). And there has been some recent interest in reviving a vagueness-indeterminacy account of OI that does not rely on that principle (see, e.g., Elson 2017). A proponent of such an account, freed from commitment to the collapsing principle, might find a “*p* is truer than *q*” rational norm to be a promising alternative to the optimizing, maximizing, and optimizing + parity rational norms. No one has yet argued that a “*p* is truer than *q*” rational norm is operative in OI cases, however, and I will not do so here. Nor will I make any further mention in Chapters 2-3 of indeterminist accounts of OI. I will once again assume, for the remainder of the dissertation, that in OI cases it is determinately false that any of the three standard value relations holds between alternatives.

Chapter 2: Epistemic Incommensurability

2.1 – Epistemic Incommensurability Cases

2.1.1 – Recap of Chapter 1

In Chapter 1, I adopted a bare-bones theory of practical rationality that I called “soft comparativism.” Soft comparativism holds that the rational justification of choice is a function of three things: (1) a comparative, quantitative fact or belief¹⁹ about the alternatives’ value; (2) a rational norm that takes the comparative, quantitative fact or belief’s propositional content as an input; and (3) any relevant deontic constraints. A “value” is anything in terms of which a value relation (better than, worse than, or equally good) may hold. The propositional content of (1) compares the alternatives with respect to whatever value is relevant to the choice. This value may itself be a composite of several values. (The value relevant to my choice of which house to buy, for example, is a composite value that includes affordability, attractiveness, neighborhood walkability, and many other constituent values.) When an agent faces a choice, (1) – (3) metaphysically ground each alternative’s deontic status (e.g., rational permissibility, obligatoriness, etc.).²⁰

Within the framework of soft comparativism, this dissertation is an attempt to answer the following question: absent any relevant deontic constraints, what does rationality require of an agent who faces a choice with two or more apparently incommensurable²¹ alternatives? This is best understood as a question about the deontic statuses of her alternatives: which alternative(s)

¹⁹ The propositional content of the relevant fact or belief may include conjunctions. Hence, no significance should be attached to my using the singular forms of “fact” and “belief.”

²⁰ I will sometimes say that something *determines* the alternatives’ deontic status. This should be read as a grounding claim.

²¹ Two alternatives are incommensurable if and only if neither is better nor worse than the other, nor are they equally good.

is (or are) rationally permissible? I aim to answer that question by providing an account of which facts, beliefs, and rational norms ground those alternatives' deontic statuses.

Cases of apparent incommensurability are either cases of ontological incommensurability ("OI") or epistemic incommensurability ("EI"). OI and EI are mutually exclusive and exhaustive. In the last chapter, I considered cases of OI. In cases of OI (if any such cases exist), the relevant comparative, quantitative fact is that incommensurable alternatives are neither better nor worse than one another, nor are they equally good. If the relevant rational norm is maximization (i.e., an alternative is rationally permissible if and only if it is *not worse than* another alternative), then any member of a set of incommensurable alternatives is rationally permissible.

On the other hand, if the relevant rational norm is optimization (i.e., an alternative is rationally permissible if and only if it is *at least as good as* each other alternative), then no member of a set of incommensurable alternatives is rationally permissible; choice between incommensurable alternatives is necessarily arational. Such choices (if optimization is the relevant rational norm) represent a breakdown of practical rationality. Parity's proponents argue that the relevant, comparative quantitative fact of incommensurable alternatives' being "on a par" saves at least some cases of OI from being necessarily irrational. If parity between alternatives is a genuine and relevant comparative, quantitative fact, then alternatives that are on a par have the same deontic status as incommensurable alternatives governed by a norm of maximization: any member of a set of on-a-par alternatives is rationally permissible, so long as there is not another alternative that is better than (and thus commensurable with) that member.²²

²² In a scenario with two alternatives, *A* and *B*, which are incommensurable but on a par, either *A* or *B* is a permissible object of choice. In a scenario with three alternatives, *C*, *D*, and *D+*, where *C* is incommensurable with but on a par with both *D* and *D+*, and *D+* is better than (and hence commensurable with) *D*, either *C* or *D+* is permissible. *D*, on the other hand, would be rationally impermissible.

In sum: in cases of OI, incommensurable alternatives either: (1) are permissible objects of choice, so long as there is not a better alternative (either because the relevant norm is maximization or because the alternatives are on a par); or (2) represent a breakdown of practical rationality (if the relevant norm is optimization).

2.1.2 – EI Cases

In this chapter, I turn to cases of EI. In EI cases, an agent's alternatives are in fact commensurable: they are either better or worse than one another, or they are equally good. However, the agent does not know which of these relations holds; she is ignorant as to her alternatives' comparative value. There is ontological commensurability; the incommensurability is merely epistemic.

EI cases, in contrast to OI cases, essentially involve agential ignorance. But more than that, they involve a specific kind of ignorance: the agent's ignorance in EI cases is (1) normative (vs. empirical); (2) not due to time constraints or inadequate deliberation; and (3) blameless.

The first two features follow from how I defined cases of “apparent incommensurability”—the topic of this dissertation—in Chapter 1. Recall that such cases are those in which the impression of incommensurability persists despite thorough reflection and a good grasp of all relevant empirical facts. It is, of course, plausible that agents are sometimes ignorant as to which alternative(s) is (or are) most valuable—and that as a result, the alternatives appear incommensurable to her—because of time constraints, or because she does not know all of the relevant empirical facts. (I might have to decide in a split second which of two strangers' lives to save²³, or a large line behind me might force me to order my drink before having fully

²³ This may strike some readers as a case of equality rather than incommensurability. Since judgments of incommensurability often arise from comparisons of qualitatively dissimilar alternatives, I invite those readers to modify the case to make the two lives qualitatively

considered each item on the menu.) But these cases are in principle resolvable: if the agent had more time to deliberate or had gathered more information, she would be able to correctly determine which alternative(s) is (or are) most valuable. (If I had more information I would realize that it is in fact better to save one stranger rather than the other—fill in the details however you like²⁴—and if I had time to fully consider the menu I would be able to accurately judge that there is a best menu item.)

My focus is instead on the seemingly irresolvable cases that have tended to get philosophers interested in the problem of incommensurability in the first place. Moreover, the effects of time constraints²⁵ and empirical ignorance²⁶ on alternatives' deontic statuses have already been well-examined elsewhere. The effects of *normative* ignorance on alternatives' deontic statuses, on the other hand, have not received the same attention. EI cases, as I will examine them here, involve ignorance only with respect to normative facts. Specifically, the agent in EI cases is ignorant with respect to which value relation (better than, worse than, or equally good) holds between her alternatives.

dissimilar in such a way that a judgment of incommensurability becomes intuitive for them. For example, perhaps one stranger is a world-class heart surgeon while the other is a world-famous concert cellist. Both lives, from a consequentialist standpoint, will bring about a more valuable state of affairs than would obtain in their absence. But they do so in such dissimilar ways that it seems impossible to judge which state of affairs (saving the surgeon or saving the cellist) is more valuable, or to judge them as equally good. They appear incommensurable.

²⁴ Perhaps one of the strangers would turn out to be a serial killer. Or, to make the case even more clear-cut, suppose that one of the strangers turns out to be a mannequin.

²⁵ There is a whole literature, for example, on Herbert Simon's (1955) proposal that when time constraints or cognitive limitations make it impossible to maximize value, it is rational to satisfice—i.e., choose an alternative that is good enough. This should not be confused with Michael Slote's notion of satisficing, according to which it is moreover sometimes rational for an agent to choose an alternative even if she judges it to be worse than another alternative (Slote 1989). The latter is premised not on agential limitations, but on the virtue of moderation.

²⁶ The literature on expected utility theory—which models the rationality of choice under empirical uncertainty—is vast.

As to the third feature, I will restrict my attention to EI cases in which the agent is a competent practical reasoner, and in which the agent's ignorance is blameless. By "blameless," I mean that her ignorance is neither intentional nor due to negligence. She has not taken pains to ignore relevant normative facts, nor has she somehow caused herself to forget normative facts that she once knew, nor does she knowingly or negligently hold contradictory beliefs about normative facts. She has made a competent, good-faith effort to learn all relevant normative facts. Once again, the purpose of these assumptions is to allow me to focus on the sorts of cases that have motivated the incommensurability literature: cases where the appearance of incommensurability persists, despite the absence of these other factors.

As I did in Chapter 1, I again set aside the question of how to determine whether a case of apparent incommensurability is a case of EI rather than OI. Answering that question—and exploring the answer's implications for practical rationality—will be the focus of Chapter 3. In this chapter, my task is to determine what rationality requires of an agent once a case of EI has been identified to be such. Nor will I argue here that EI cases exist—i.e., that not all cases of apparent incommensurability are OI cases. In this chapter I will take the existence of EI cases for granted, and focus instead on the question of what practical rationality requires of agents facing such cases.

Surprisingly, this question has been almost entirely neglected by philosophers working on value incommensurability. There are, I think, two reasons for this neglect. The first is that for a long time, the value incommensurability literature focused on arguments for and against the existence of OI and EI cases (with the overwhelming majority focusing on OI cases).²⁷ And

²⁷ See, e.g., many of the essays in Chang 1997a. Of these, only one—Regan 1997—presents a sustained defense of the existence of EI cases. For a useful taxonomy of arguments for the existence of OI cases, see Chang 1997b, § III.

while the OI literature has shifted more recently to include discussions of alternatives' rational permissibility,²⁸ the EI literature—which was much sparser to begin with—has not followed suit.

The second reason stems from the fact that EI cases involve agential ignorance. The effects of ignorance or uncertainty²⁹ on the rationality of choice have been extensively studied, but mostly in the context of expected utility theory. Expected utility theory, however, is concerned with *empirical* uncertainty (e.g., whether or not it will rain later today, or which number will result from a die roll). Agents are modeled as being uncertain as to which of several possible states of the world obtains. *Value*—utility—is grounded in agential preferences, and the idealized agent of expected utility theory has complete preferences. There is a kind of normative ignorance, but it is solely a result of empirical ignorance. Because it is uncertain which state obtains, alternatives' (or acts') deontic statuses are determined by *expected* utility, rather than actual utility. Within each state, though, each act has a known utility value.³⁰

Shockingly little has been written about the effects of *normative* uncertainty on practical deliberation. A handful of philosophers are working to rectify this oversight, mostly by adapting existing tools from formal decision theory to choices that involve agential uncertainty about normative facts.³¹ This literature, while promising, has not yet addressed the issue of epistemic incommensurability.

²⁸ The primary driver for this has been an ongoing fleshing-out of the theory of parity by Ruth Chang (2002a, Ch. 5-6; 2002b; 2005; 2016b), Wlodek Rabinowicz (2008; 2009a; 2009b; 2012), and Joshua Gert (2004; 2015).

²⁹ I use the term “uncertainty” rather than “ignorance” here because the term “ignorance” as used by decision theorists has a specific, narrower meaning (see below §§ 2.2-2.3).

³⁰ Depending on how acts are characterized, it might be more accurate to say that each act *yields* a utility value. According to Leonard Savage's (1954, § 2.5) classic account, for example, an act just is a function from states to consequences (now typically called outcomes), which are measured in utility.

³¹ See, e.g., Sepielli 2009, 2013, 2014; Ross 2006, and Lockhart 2000. Ross's and Lockhart's work focuses specifically on the issue of uncertainty as to which normative ethical theory (e.g.,

In short, there is a growing literature on rational permissibility in OI cases (at least with respect to parity), and a rich literature on agential uncertainty, but virtually no literature on how agential uncertainty affects rational permissibility in EI cases. This chapter is my attempt to fill that void. It extends the research agenda of the small-but-growing literature on normative uncertainty to the phenomenon of epistemic incommensurability.

What does rationality require of an agent who faces commensurable alternatives but remains, despite adequate or even exhaustive efforts, ignorant as to the alternatives' relative value? Put differently: in EI cases, what grounds alternatives' deontic statuses? In the absence of deontic constraints, alternatives' deontic statuses are jointly determined by (1) a fact or belief and (2) a rational norm. My inquiry will focus on rational norms, since different rational norms take different propositional content as inputs. A norm of maximizing expected value, for example, takes (*inter alia*) probabilities of possible states of the world as inputs, while many other norms (e.g., maximizing actual value; maximin; maximax) do not. Hence, answering the question of which rational norm is operative in EI cases will also answer the question of which fact or belief grounds alternatives' deontic statuses. The rational norm and fact or belief operate as a pair, jointly determining alternatives' deontic statuses.

As for the rational norm, the most obvious answer is that a maximizing or optimizing norm determines the alternatives' deontic statuses, just as it does in OI cases and in ordinary choices that don't involve any apparent incommensurability. But figuring out precisely *what* to maximize or optimize is less straightforward in EI cases. In section 2.2, I consider the rational norm of maximizing *actual* value. I distinguish objective and subjective deontic statuses, and

Kantianism, utilitarianism, egoism) is correct. Sepielli 2009 makes the important contribution of generalizing their approach to include other cases of normative uncertainty. For a criticism of this approach, see Harman 2015. Thanks to Russ Shafer-Landau for introducing me to this literature.

argue that the norm of maximizing actual value grounds objective—but not subjective—deontic statuses. In section 2.3, I ask whether a norm of maximizing *expected* value may ground subjective deontic statuses. I argue that because EI cases are better understood as decisions under ignorance than decisions under risk, the norm of maximizing expected value cannot by itself ground subjective deontic statuses. I then consider whether it may do so with an added assumption: the principle of insufficient reason. I consider several objections to the principle of insufficient reason, and argue that one—the Equiprobability Objection—is fatal. In section 2.4, I argue that Maximin—a norm of maximizing *minimum possible value*—grounds subjective deontic statuses in EI cases. In the course of doing so, I reject three other maximizing norms: Maximax (a norm of maximizing *maximum possible value*), the MAPV norm (a norm of maximizing *average possible value*), and Medianmax (a norm of maximizing *median possible value*).

2.2 – Maximizing Actual Value

My project in this chapter is to provide an account of what grounds alternatives' deontic statuses in EI cases. A natural starting point is to ask whether EI cases are governed by the same rational norm as OI cases. If this were the case, alternatives' deontic statuses would be grounded in a norm of maximizing or optimizing actual value. (Call this the "MAV norm.") By "actual value," I mean the value that each alternative in fact has, regardless of what the agent believes about its value. Since the alternatives in EI cases are commensurable, we need not specify whether the relevant norm is maximization or optimization. If a set of alternatives is fully commensurable, the set of alternatives that maximizes is identical to the set of alternatives that optimizes. Hence, both maximization and optimization norms will render the same set of prescriptions. In EI cases, if the MAV norm is the operative rational norm, an alternative is

rationaly permissible (absent deontic constraints) if and only if it is at least as good as (equivalently: not worse than) each other alternative. (For convenience, I will drop the term “optimization” and use only the more popular “maximization” for the rest of the chapter.)

This account offers the attractive features of consistency and parsimony: EI cases and OI cases are governed by a single norm. And a norm that determines alternatives’ rational permissibility solely on the basis of their comparative value, with no reference to the agent’s beliefs, fits with the model of practical rationality that dominates current work on value incommensurability. Ruth Chang, for example, adheres to a version of comparativism in which alternatives’ deontic statuses are determined solely by “comparative facts about the evaluative merits of the options” (2016a, 214; see also 2004a; 2004b).

Suppose that an agent faces a two-alternative EI case. She must choose between alternative *A* and alternative *B*. She knows all of the relevant empirical facts, and has deliberated so thoroughly that, absent a fresh epiphany that seems increasingly unlikely, any further deliberation will be redundant. Despite her best efforts, the cases remain apparently incommensurable to her: she cannot confidently judge that one is better than the other, nor that they are equally good. Yet there is a fact of the matter as to *A* and *B*’s relative value: one of these comparative relations holds between *A* and *B*.

Under the MAV norm, the agent’s ignorance about normative facts has no affect on the rational permissibility of her alternatives. The rational permissibility of her alternatives is straightforwardly determined by the fact of *A* and *B*’s comparative value, together with a maximizing norm. There are three possibilities: (1) one alternative is permissible (because it is better than the other), and she chooses that alternative; (2) one alternative is permissible (because it is better than the other), but she chooses the impermissible alternative; (3) either alternative is

permissible (because they are equal), and so she chooses a permissible alternative. Under any of these three possibilities, the agent chooses a rationally permissible alternative, if at all, only by accident (or luck). Similarly, if she chooses a rationally impermissible alternative, she does so coincidentally. This is true even in the equality case. She has no epistemic access to the fact that grounds her alternatives' rational permissibility; as a matter of luck, she cannot help but choose a rationally permissible alternative.

In possibilities (1) and (3), the agent chooses the “right” alternative. Conversely, she chooses the “wrong” alternative in (2). But it does not follow that she chose a rationally permissible alternative unless we further assume that rightness entails rational permissibility. The premise that rightness entails rational permissibility, when taken together with the MAV norm, suggests the following picture. Suppose that *A* is in fact better than *B*. If the agent would act rationally (i.e., choose a rationally permissible alternative), she must choose *A*. Yet if she does choose *A*, there is another sense in which it seems inappropriate to say on that basis alone that she acted rationally (despite the fact that she chose a rationally permissible alternative). She chose the right alternative—the value-maximizing alternative—but only as a matter of luck. Acknowledging that she chose the *right alternative* does not automatically resolve the question of whether she *acted rightly*—i.e., chose a rationally permissible alternative.

The confusion here arises from the fact that there are two senses of the term “rightness” that must be distinguished: objective and subjective.³² *Objective rightness* consists in something being right in light of the agent's actual circumstances. This is the sense of rightness that is operative when we say that the agent chose the right alternative. *Subjective rightness*, by

³² My discussion of objective and subjective rightness draws heavily on Smith 2010b, which provides the clearest explication of this distinction that I've seen.

contrast, consists in something being right relative to the agent's beliefs.³³ This is the sense of rightness that is operative when we discuss whether the agent acted rightly. Each sense of rightness corresponds with a different set of deontic statuses. Objective rightness corresponds with objective obligatoriness, permissibility, etc., while subjective rightness corresponds with subjective obligatoriness, permissibility, etc.

The distinction between objective and subjective rightness is primarily used by moral philosophers.³⁴ In that context it initially arose as a way of making sense of our tendency, when employing consequentialist theories that take far-flung and impossible-to-determine empirical consequences as relevant to the normative question of whether an action is morally right,³⁵ to judge that an agent simultaneously acted rightly and wrongly. If I save toddler Hitler from drowning in a pool, for example, this sort of consequentialist theory would take the action as objectively wrong: my chosen alternative was a state of affairs that included the Holocaust and the mass genocide of around 6 million Jews. Yet the action was subjectively right: given my ignorance about young Adolf's trajectory, it was (subjectively) right to save the child from drowning.

³³ It is a point of contention whether "beliefs" here means (1) the agent's actual beliefs or (2) some sufficiently laundered version of those beliefs (e.g., what it would be reasonable for the agent to believe, or what an ideal agent in her circumstances would believe). See Smith 2010a (arguing that subjective rightness is grounded in the agent's actual beliefs, rather than in the beliefs that it would be reasonable for her to have).

³⁴ See Smith 2010b, n.4 for an extensive catalog of contemporary moral theorists who employ this distinction—though, as Smith notes, the distinction is not universally accepted.

³⁵ See, e.g., G.E. Moore's *Principia Ethica*:

In order to shew [*sic*] that any action is a duty, it is necessary to know ... all the events which will be in any way affected by our action throughout an infinite future. ... But it is obvious that our causal knowledge alone is far too incomplete for us ever to assure ourselves of this result. Accordingly it follows that we never have any reason to suppose that an action is our duty: we can never be sure that any action will produce the greatest value possible. (1903, § 91)

The objective-subjective rightness distinction is not limited to moral consequentialism. Holly Smith (2010b, 66-67) argues convincingly that it applies equally to deontological moral theories. Nor is it limited to the moral domain. In fact, it is hard to imagine a domain of practical reason where the distinction would not be meaningful. In prudential reasoning, for example, a high-paying job may be both the subjectively right alternative and the objectively wrong alternative—say, because it turns out that (unbeknownst to me) the boss is a tyrant or the employer will go out of business within the year. Similarly, in epistemic reasoning we may distinguish what is (objectively) right to believe in light of an agent's evidence from what is (subjectively) right to believe in light of the agent's beliefs about her evidence.

The premise that rightness entails permissibility, then, may be split into two premises: one concerning objective rightness and permissibility, the other concerning subjective rightness and permissibility. The MAV norm plausibly determines which alternative is objectively right, and *a fortiori*, which alternative is objectively rationally permissible. Absent deontic constraints, an alternative is objectively right, and hence objectively rationally permissible, if and only if it maximizes actual value.

This is true for EI cases for all of the same reasons that it was true of OI cases in Chapter 1. In OI cases there was no need to separately consider subjective rightness, because the agent's beliefs about her circumstances matched reality. More precisely, her belief about the alternatives' comparative value and the fact of their comparative value had identical propositional content: the alternatives were incommensurable.

But EI cases, unlike OI cases, involve ignorance. The comparative fact that determines, together with the MAV norm, which alternative is objectively rationally permissible is not known to the agent. If the MAV norm were the relevant norm for *subjective* rightness as well,

then the agent's ignorance would be irrelevant to a determination of subjective rightness—the same comparative fact would determine subjective rightness, despite her ignorance of that fact.

However, it makes no sense to apply the MAV norm to subjective rightness in EI cases. Subjective rightness by definition takes the agent's beliefs about her circumstances, rather than the actual facts of those circumstances, as determinative. It does so in order to fulfill at least two critical functions of a theory of practical rationality that cannot be fulfilled by objective rightness: (1) providing guidance for agents' decisions; and (2) assigning praise or blameworthiness. A theory of practical rationality that conditions subjective rational permissibility on facts of which the agent is (non-culpably) ignorant fails to fulfill both of these functions. It can provide no guidance to an agent facing an EI case, since it grounds rightness in a fact that she does not know, and of which she remains intractably ignorant. And assigning blame to an earnest, competent agent for failing to choose the right alternative as a matter of bad luck runs afoul of any plausible account of blameworthiness.

So while the MAV norm plausibly grounds the objective deontic statuses of alternatives in EI cases, it does not ground alternatives' subjective deontic statuses. Yet subjective deontic statuses are a crucial component of practical rationality. They both serve as a basis for guiding agents' decisions and ground ascriptions of blameworthiness/praiseworthiness. An account of practical rationality in EI cases would be incomplete without them. Because the MAV norm does not ground subjective deontic statuses, we must look for some other rational norm that does. In the next section, I will consider the most obvious contender: a norm of maximizing expected, rather than actual, value.

2.3 – Maximizing Expected Value and the Principle of Insufficient Reason

2.3.1 – The MEV Norm

According to the norm of maximizing expected value (the “MEV norm”), an alternative is rationally permissible (absent deontic constraints) if and only if it maximizes expected value—i.e., its expected value is at least as good as (equivalently: not worse than) that of each other alternative. An alternative’s expected value is a probability-weighted average of its possible values. It is equal to the sum of the value that the alternative would have in each of several possible states of the world multiplied in each case by the probability of that state’s obtaining:

$$\begin{aligned} \text{Expected value of alternative } A = & (\text{value of } A \text{ under state 1})(\text{probability of state 1}) \\ & + (\text{value of } A \text{ under state 2})(\text{probability of state 2}) + \dots + (\text{value of } A \text{ under state } \\ & n)(\text{probability of state } n) \end{aligned}$$

The possible states of the world are states of affairs. For brevity, I will sometimes refer to them simply as “states.” The states that pertain to the MEV norm must be mutually exclusive and exhaustive, so that they form a partition and their probabilities sum to 1.³⁶

I have formulated expected value in terms of the expected value of *alternatives* because this fits with my general presentation of soft comparativism, according to which the rationality of choice is a function of alternatives’ value. I have chosen to model the rationality of choice as a function of alternatives’ value because that is the dominant approach in discussions of value relations and incommensurability. Nonetheless, I acknowledge that this formulation may seem idiosyncratic to decision theorists. In that discipline, the convention is to model *acts* (which are typically taken to be mutually exclusive and exhaustive *alternatives*), which have a different *outcome* under each state,

³⁶ The fact that the probabilities of mutually exclusive and exhaustive states sum to 1 follows straightforwardly from the basic axioms of the probability calculus. Rather than rehearse this proof here, I will refer curious readers to Resnik 1987, § 3-2, Theorem 1.

and to formulate the MEV norm as maximizing expected values of outcomes—not acts or alternatives. Fortunately, none of my substantive claims hinges on the fact that I characterize the alternatives, rather than their outcomes under states, as bearers of value.

While the MAV norm necessarily pertains to objective rightness, the MEV norm may pertain to either objective or subjective rightness—depending on whether the probabilities it takes as inputs are objective probabilities or (potentially mistaken or incomplete) subjective probabilities. By “subjective probabilities,” I mean probabilities that are grounded in the agent’s beliefs, rather than in her actual circumstances. A subjective probability is not necessarily a belief whose propositional content has an explicit probability—e.g., a belief that “there is a $2/7$ probability that I will roll a 5,” or a belief that “there is a $1/3$ probability that *A* is better than *B*.” It may instead be some other feature of her doxastic attitudes. To the extent that I can, I will remain neutral on this question of how subjective probabilities are best characterized. For my purposes, the crucial feature of subjective probabilities is that they are grounded in the agent’s beliefs.

Suppose, for example, that an agent is deciding whether to place a bet on the outcome of a six-sided die roll. It is a fair die, but the agent incorrectly believes it is weighted. In that case, the MEV norm plausibly determines the objectively right bet by taking as its input a probability of $1/6$ for each face, while also plausibly determining the subjectively right bet by taking as its input the agent’s subjective probability for each face (say, $1/7$ for five of the faces, and $2/7$ for the sixth face).

The MEV norm is the natural starting point in our search to determine which rational norm grounds alternatives’ subjective deontic statuses in EI cases. It is the most popular and influential rational norm for modeling decisions under uncertainty, and EI cases are decisions

under uncertainty: the agent is uncertain as to her alternatives' comparative value. There is also some precedent for regarding the MEV norm as the ground of subjective rightness in choices that involve normative uncertainty. Lockhart 2000 (see especially Ch. 2) argues that a version of the MEV norm is operative in choices that involve moral uncertainty, and Sepielli 2009 argues that a more generalized version of the MEV norm is operative in all decisions that involve normative uncertainty.

Moreover, the MAV norm may itself be reformulated in terms of the MEV norm. Here is how this would work in a two-alternative EI case where A is in fact better than B . Let the possible states of the world consist of the following set: $\{A$ is better than B , A is worse than B , A is equally as good as $B\}$. The probability of the state " A is better than B " is 1; each other state has a probability of 0. Hence, the expected value of A —1 times the actual value of A —is higher than the expected value of B —1 times the actual value of B . So long as we specify the states and probabilities in this way, an alternative will maximize expected value if and only if it maximizes actual value. The fact that the norm governing *objective* deontic statuses may be reformulated in MEV terms provides another reason for starting our search for the determinant of *subjective* deontic statuses with the MEV norm. If both objective and subjective deontic statuses may be plausibly grounded in a single norm, then considerations of parsimony and consistency provide an additional reason (beyond its plausibility) to accept that norm as genuine.

Unfortunately, applying the MEV norm as a determinant of *subjective* deontic statuses in EI cases is less straightforward and more controversial. Before we can apply the MEV norm, we must first settle on plausible descriptions of both the relevant states and the relevant probabilities. Indeed, strictly speaking, we do not *have* a rational norm at all until we have an

account of the relevant states and probabilities.³⁷ The set of states in a two-alternative EI case may be characterized in the same way for subjective deontic statuses as for objective deontic statuses: $\{A \text{ is better than } B, A \text{ is worse than } B, A \text{ is equally as good as } B\}$. This is admittedly an unusual way to characterize states, at least from the standpoint of formal decision theory. In that context, states typically represent empirical uncertainty. But in that context there is no normative uncertainty, because the agent is typically assumed to have complete preferences. In EI cases, in contrast, there is *only* normative uncertainty. The agent knows all of the relevant empirical facts. If the MEV norm is to apply to EI cases at all, then, the states must be partitioned in such a way as to represent the agent's normative uncertainty about her alternatives' relative values. Although this approach to state descriptions is unorthodox, I do not think it is problematic. It does not, for example, include states that are causally dependent on the agent's choice.³⁸ We have, then, a plausible description of the relevant states.³⁹

³⁷ Thanks to Peter Vranas for clarifying this point for me.

³⁸ The fact of a state's being causally dependent on which alternative is chosen does not necessarily mean that the set of state descriptions is incorrect or illegitimate, but it does complicate things in a couple of ways. First, it follows that dominance reasoning is inapplicable to the decision, because dominance reasoning may yield an obviously wrong prescription (see Reznik 1987, § 1-2a; Peterson 2009, 20-21). Second, it follows that the MEV norm does not take *the states' probabilities* as inputs. Instead, it takes *their probabilities conditioned on the alternative's being chosen* as inputs (Jeffrey 1983). In other words:

$$\text{Expected value of alternative } A = (\text{value of } A \text{ under state 1})(\text{probability of state 1 conditional on } A) + (\text{value of } A \text{ under state 2})(\text{probability of state 2 conditional on } A) + \dots + (\text{value of } A \text{ under state } n)(\text{probability of state } n \text{ conditional on } A)$$

Call this "Jeffrey's formulation" of the MEV norm, and call the formulation of the MEV norm that I initially articulated "Savage's formulation" (after Savage 1954). When the states and alternatives are causally independent, Jeffrey's formulation reduces to Savage's formulation. (If the probabilities are independent, then the probability of state ϕ conditioned on alternative A just is the probability of state ϕ .) Since the states in EI cases as I have described them are causally independent of the alternatives, we may use Savage's more straightforward formulation.

³⁹ More specifically, I take the set of states to be a mutually exclusive and exhaustive list of all of the possible ways that the three standard value relations might hold between the alternatives. Thus, for example, a three-alternative EI case of complete apparent incommensurability would have fifteen states: $\{A \text{ is better than } B \text{ and } B \text{ is better than } C, B \text{ is better than } C \text{ and } C \text{ is better than } A, C \text{ is better than } A \text{ and } A \text{ is better than } B, A \text{ is better than } B \text{ and } C \text{ is better than } A, B \text{ is better than } C \text{ and } A \text{ is better than } B, C \text{ is better than } A \text{ and } B \text{ is better than } C, A \text{ is better than } B \text{ and } C \text{ is better than } B, B \text{ is better than } C \text{ and } A \text{ is better than } C, C \text{ is better than } A \text{ and } B \text{ is better than } A, A \text{ is better than } C \text{ and } B \text{ is better than } A, B \text{ is better than } A \text{ and } C \text{ is better than } B, C \text{ is better than } B \text{ and } A \text{ is better than } C, A \text{ is better than } C \text{ and } B \text{ is better than } C, B \text{ is better than } A \text{ and } C \text{ is better than } A\}$

The difficulty in applying the MEV norm to subjective rightness in EI cases lies in trying to settle on a plausible characterization of the relevant subjective probabilities. The chief question, in fact, is whether there are any relevant probabilities at all. For objective rightness, the relevant probabilities are each state's actual probability. But for subjective rightness, the relevant probabilities are subjective. They are grounded in the agent's doxastic attitudes with respect to each state. What is the correct way to characterize an agent's doxastic attitudes, in a two-alternative EI case, with respect to the states $\{A \text{ is better than } B, A \text{ is worse than } B, A \text{ is equally as good as } B\}$? She has thoroughly deliberated, but remains flummoxed as to which value relation holds between A and B . It is not even that the case is a "close call." Rather, A and B appear to her to be incommensurable. Yet she is confident that despite appearances, her alternatives are in fact commensurable. One might be tempted to say that she views each state as equiprobable, in which case each state has a probability of $1/3$. (More on equiprobability below.) But this does not quite seem to be an accurate description of her doxastic attitudes. Rather, she feels—in most cases because her alternatives are so qualitatively dissimilar—that she cannot even begin to estimate which state is most likely to be actual, let alone that they are equally likely. All of her deliberative efforts lead her back to the perception of apparent incommensurability.

than A , C is better than A and A is better than B , A is better than C and C is better than B , B is better than A and A is better than C , C is better than B and B is better than A , A is better than B and B is equally as good as C , B is better than A and A is equally as good as C , A is equally as good as B and A is better than C , A is better than B and B is equally as good as C , C is better than A and A is equally as good as B , A is equally as good as C and A is better than B , B is better than A and A is equally as good as C , C is better than A and A is equally as good as B , B is equally as good as C and B is better than A . I assume that the value relations are transitive—so that, for example, when I say of the first state that " A is better than B and B is better than C ," it follows that A is better than C . Thanks to Peter Vranas for pushing me to say more here about my general method of defining states.

She is facing what decision theorists call a decision under ignorance, rather than a decision under risk. In decisions under risk, an agent knows (or can at least estimate) the probability of each possible state. In decisions under ignorance, however, the probabilities of possible states are either unknown or nonexistent. In other words, there are no subjective probabilities. This leads to the following quandary: the MEV norm takes probabilities as inputs. But EI cases are decisions under ignorance. Hence, in EI cases, there are no relevant subjective probabilities for the MEV norm to take as inputs. Accordingly, the MEV norm cannot ground alternatives' subjective deontic statuses. Or, at least, it cannot do so without our taking on an ancillary assumption.

2.3.2 – The Principle of Insufficient Reason

The principle of insufficient reason is an ancillary assumption that allows the MEV norm to apply to decisions under ignorance. As typically stated, it holds that where there is insufficient reason to regard one state as more probable than another, the agent should view the states as equally probable.⁴⁰ When formulated in this way, the principle of insufficient reason is a rational constraint on the agent's beliefs. It informs the MEV norm only if we make the further assumption that the MEV norm, when determining subjective rightness, takes beliefs that it would be reasonable (i.e., rationally permissible) for an agent to have, rather than her actual beliefs, as an input. This assumption is contested (see Smith 2010a). It is also irrelevant to my discussion of EI cases, because I have stipulated that in EI cases the agent's actual beliefs are reasonable.

We can drop the assumption by reformulating the principle as follows:

⁴⁰ See, e.g., Sinn 1980 (“The Principle of Insufficient Reason ... says that if there is no reason to believe that out of a set of possible, mutually exclusive events no one event is more likely to occur than any other, then one should assume that all events are equally probable.”).

principle of insufficient reason: If an agent's beliefs provide no reason to suppose that one state is more probable than another, then the MEV norm⁴¹ treats the states as equally probable.

“Beliefs” here may be interpreted to mean either the agent's actual beliefs or the beliefs that it would be reasonable for the agent to have, given her circumstances. At the beginning of the chapter I defined EI cases in such a way that these two interpretations have the same extension. The agent, having competently and thoroughly deliberated, holds the beliefs that it would be reasonable for her to have in the circumstances. Nevertheless, the alternatives appear incommensurable to her. If we partition the possible states of the world according to which value relation holds between the apparently incommensurable alternatives, her beliefs provide no reason to suppose that one state is more probable than another.

In a two-alternative EI case, it follows from the principle of insufficient reason that the MEV norm takes as inputs a 1/3 probability that *A* is better than *B*, a 1/3 probability that *A* is worse than *B*, and a 1/3 probability that *A* is equally as good as *B*. Before we can apply the MEV norm to a two-alternative EI case, though, we further need to be able to numerically represent the alternatives' value under each state. Here my unorthodox method of partitioning the states forces us into an awkward position, as illustrated by the following decision matrix⁴²:

⁴¹ Though I will discuss the principle of insufficient reason only in combination with the MEV norm, it may in theory be combined with any rational norm that takes probabilities as inputs. Here is a more cumbersome reformulation that is neutral between rational norms: if (1) an agent's beliefs provide no reason to suppose that one state is more probable than another, and (2) the rational norm that grounds her alternatives' subjective deontic statuses takes those states' probabilities as an input, then (3) that rational norm treats the states as equally probable.

⁴² I follow the usual decision matrix conventions. The leftmost column lists alternatives; the topmost row depicts partitioned states; the remaining squares list the values that result from an alternative-state pair. The only anomaly is that I take the alternatives, rather than their outcomes under states, to be the bearers of the value relevant to the choice. But as discussed earlier, I do not think any of my substantive claims depends on this.

Table 2.1

	<i>A</i> is better than <i>B</i>	<i>A</i> is worse than <i>B</i>	<i>A</i> is equally as good as <i>B</i>
<i>A</i>	2	1	?
<i>B</i>	1	2	?

The issue here is that the MEV norm applies an alternative's values *across* states. In the better-than and worse-than states, the numbers simply represent the fact that *A* and *B*'s comparative values reverse.⁴³ But what numeric value ought we apply to the equality state? *A* and *B* ought have the same value as one another, but attempts to assign any particular number seem problematic. An assignment of "1" to both alternatives wrongly suggests that the alternatives, when equal to one another, are less valuable than when they are the better alternative. A similar problem applies if we assign "2." We could split the difference and assign a value of "1.5," but doing so seems ad hoc. (Why not "1.6"?)

Fortunately, it does not matter which number we choose to represent the alternatives' values in the equality state. The choice of number does not affect whether one alternative's expected value is higher than the other's, because it is effectively a constant that is added to both:

$$\begin{aligned} \text{Expected value of } A &= (1/3)(2) + (1/3)(1) + (1/3)(k) \\ \text{Expected value of } B &= (1/3)(1) + (1/3)(2) + (1/3)(k) \end{aligned}$$

⁴³ In truth, the values in the first two states are a bit more contentious than the above description suggests. To arrive at them, I must make three assumptions: (1) that the states' minima are the same; (2) that the states' maxima are the same; and (3) that the interval between *A*'s and *B*'s values is the same in both states. (Thanks to Peter Vranas for stating this so clearly for me.) (1) and (2) I justify on the basis that we have no grounds to suppose that the minima and maxima differ. I cannot imagine why, for example, if *A* is worse than *B* (i.e., if state 2 obtains), it would be lower in value than *B* would be if *B* were worse than *A* (i.e., if state 1 obtains). And once we assume (1) and (2), (3) follows. As for the particular interval between *A* and *B*, it is true that I cannot help but take the values in Table 2.1 to represent an interval scale. (This is because, as we'll see, the MEV norm requires interval-measurable values.) But in a case of complete apparent incommensurability, as in Table 2.1, the particular interval that I represent will not change the MEV norm's prescription. And for Maximin—the norm that I ultimately endorse as operative in EI cases—we need not assume interval-measurability at all. Ordinal values are sufficient.

The alternatives' expected values are equal: $1 + 1/3(k)$. According to the MEV norm, this means that A and B are both subjectively rationally permissible objects of choice. If we accept both the principle of insufficient reason and the MEV norm, then, the agent in a two-alternative EI case may choose either alternative.

Should we accept the principle of insufficient reason as applicable to EI cases? In the remainder of section 2, I will consider what I take to be the three strongest objections to the proposal that the principle of insufficient reason applies (alongside the MEV norm) to EI cases: the Multiple Partitions Objection, the Interval Scale Objection, and the Equiprobability Objection. I will argue that because of the particular features of EI cases, the Multiple Partitions Objection is not fatal. I will then argue that the Interval Scale Objection, while not fatal, nonetheless requires us to find another norm that is capable of operating in EI cases in which the relevant value admits of merely ordinal (and not cardinal) comparisons. Finally, I will argue that the Equiprobability Objection makes the principle of insufficient reason implausible even in EI cases whose relevant value admits of cardinal comparisons. Because there are no subjective probabilities in EI cases, the MEV norm cannot operate without an ancillary assumption like the principle of insufficient reason. Because we should reject the principle of insufficient reason and because there is no other plausible ancillary assumption that can take its place, we must also reject the MEV norm as operative in EI cases. We must thus search for another rational norm that may plausibly ground subjective rightness in EI cases. This will be my task in section 3.

2.3.3 – The Multiple Partitions Objection

One potential problem with the principle of insufficient reason is that when combined with the MEV norm, the resulting prescriptions are notoriously sensitive to how states are partitioned. But for some choices, it is implausible to suppose that there is exactly one correct

way to partition the states. If the principle of insufficient reason and the MEV norm jointly ground subjective rightness, then there may be choices in which two equally plausible partition methods result in inconsistent deontic statuses. Call this the “Multiple Partitions Objection.”

Suppose, for example, that I am trying to decide whether to water my garden this evening. The value relevant to my choice is garden productivity. If I water and it also rains overnight, my tomatoes will split. On the other hand, if I don’t water and it doesn’t rain, the seeds I planted yesterday will dry out and fail to germinate. I need the garden watered one way or another; it makes no difference to me whether by rain or by hose. For clarity, we may summarize all of this into the following decision matrix:

Table 2.2

	rain	no rain
water	0	1
don’t water	1	0

Here, the possible states are partitioned as “rain” and “no rain.” The expected value of either alternative (“water” or “don’t water”) is $\frac{1}{2}$. So it would seem that under the principle of insufficient reason and the MEV norm, either alternative is permissible.

But suppose instead that we partition the states as “heavy rain,” “moderate rain,” and “no rain.” If I water, moderate rain will still be enough to split my delicate heirloom tomatoes:

Table 2.3

	Heavy rain	Moderate rain	No rain
Water	0	0	1
Don’t water	1	1	0

The expected value of watering is $\frac{1}{3}$, and the expected value of not watering is $\frac{2}{3}$. Now it seems that under the principle of insufficient reason and the MEV norm, watering is rationally impermissible and not watering is obligatory.

Now suppose that we further subdivide “moderate rain” into “moderate rain” and “light rain.” Light rain may not be enough to keep my seedlings moist. And so:

Table 2.4

	Heavy rain	Moderate rain	Light rain	No rain
Water	0	0	1	1
Don't water	1	1	0	0

Now the expected values are both $\frac{1}{2}$ again, and either alternative is permissible.

The problem here is that for many sorts of choices, as in the garden-watering example, it is implausible to suppose that there is exactly one correct way to partition the states. But unless there is a single correct way to partition the states, the principle of insufficient reason may yield contradictory deontic statuses. In my example, it is both permissible and impermissible to water my garden.

In fact, the problem is worse than my garden-watering example suggests. That example, and analogous choices, might be resolvable by appealing to the notion of coarse-grained versus fine-grained partition methods. There are infinite levels of fine-grainedness,⁴⁴ but perhaps there is a principled way to specify exactly how fine- or coarse-grainedly to partition the states for the purposes of grounding subjective deontic statuses. Even so, there may be other choices that run up against a problem known as the Bertrand paradox.

Here is an example from van Fraassen (1989, 303-304) that illustrates the paradox:

A precision tool factory produces iron cubes with edge length ≤ 2 cm. What is the probability that a cube has length ≤ 1 cm, given that it was produced by that factory?

According to a “naïve application” of the principle of insufficient reason, there is a $\frac{1}{2}$ probability that the cube has an edge length between 0 cm and 1 cm (i.e., $0 < l \leq 1$), and a $\frac{1}{2}$ probability that

⁴⁴ Suppose we specify some maximal amount of rain that might fall overnight. We might subdivide the states according to half that quantity of rain, or half of *that* quantity of rain, and so on.

the cube has an edge length between 1 cm and 2 cm (i.e., $1 < l \leq 2$). But suppose we ask instead whether the cube has a side area that is between 0 cm and 1 cm. The two questions are logically equivalent: a cube has a side area between 0 cm and 1 cm if and only if its edge length is between 0 cm and 1 cm. Instead of partitioning into two states, however, the side-area question results in a partition with four states: $\{0 < a \leq 1, 1 < a \leq 2, 2 < a \leq 3, 3 < a \leq 4\}$. Now, according to the principle of insufficient reason, the probability of the state with edge length ≤ 1 cm is $\frac{1}{4}$. But we may also ask the logically equivalent question of whether the cube has a volume ≤ 1 cm. Because the maximum volume is 8 cm, we now have eight partitions, and the probability of the state with an edge length ≤ 1 cm is $\frac{1}{8}$. So: three logically equivalent descriptions (length, area, or volume ≤ 1 cm) yield three different partitions, with three different probabilities that the edge length is ≤ 1 cm.

The fact that the principle of insufficient reason, as applied to three logically equivalent propositions, yields three logically inconsistent probabilities is itself problematic. And it is not hard to see how the Bertrand paradox might further result in the principle of insufficient reason and the MEV norm grounding inconsistent subjective deontic statuses. Suppose that I am a jewelry maker, and I need a metal cube with edge length ≤ 1 cm. For some reason—never mind why—my metal supplier offers me a choice between (1) a 50:50 chance of either a cube with edges between 0 cm and 1 cm, or a cube with edges between 1 cm and 2 cm; and (2) whatever cube will next come off the assembly line at van Fraassen's factory. For simplicity, let's numerically represent the value of a useful cube as 1, and the value of a too-big cube as 0. The first alternative has an expected value of $\frac{1}{2}$. The second alternative has either an expected value of $\frac{1}{2}$, $\frac{1}{4}$, or $\frac{1}{8}$, depending on which of van Fraassen's three partitions we use. Under the edge-

length partition, either alternative is (subjectively) rationally permissible. But under the side-area and cube-volume partitions, the factory cube is (subjectively) rationally impermissible.

The Multiple Partitions Objection is a serious problem for the principle of insufficient reason if we take it as a generally applicable principle of practical reason. Here, however, our task is more limited: to determine whether the principle of insufficient reason plausibly applies to determinations of subjective rightness in EI cases. And in EI cases, it is actually plausible that there is exactly one correct way to partition the states. The agent is ignorant only as to which value relation (better than, worse than, or equally good) holds between her alternatives. The states must be partitioned in a way that reflects this ignorance. Here are all of the logically possible ways of doing so:

1. { A is better than B , A is worse than B , A is equally as good as B }
2. { A is at least as good as B , A is worse than B }
3. { B is at least as good as A , B is worse than A }
4. { A is either better or worse than (but not equally as good as) B , A is equally as good as B }
5. { A is either better than, worse than, or equally as good as B }

The MEV norm takes numerical representations of the alternatives' values under each state as an input. Accordingly, a partition method that does not permit sufficient numerical representation of the alternatives' values for the MEV norm to ground deontic statuses (even when taken together with the principle of insufficient reason) is worse than a partition method that does. On these grounds we can rule out partition method #4 as a worse way of partitioning the states than the others. There is no way to numerically represent A as being worse than or better than, but not equally as good as, B . As a result, if we use partition method #4, then the MEV norm cannot determine the alternatives' subjective deontic statuses.

Partition methods #2, #3, and #5 are also demonstrably worse methods than #1. Consider method #2. Representing the values under the state " A is worse than B " is simple enough: assign

a lower number (say, 0) to B and a higher number (say, 1) to A . For the state “ A is at least as good as B ,” however, we may *either* represent A and B as equally good (as in Table 2.5) or represent A as better than B (as in Table 2.6):

Table 2.5

	A is at least as good as B	A is worse than B
A	k	0
B	k	1

Table 2.6

	A is at least as good as B	A is worse than B
A	1	0
B	0	1

(As before, the effect of adding a state where the alternatives are equally good is equivalent to adding a constant, k , to the expected values under the other states.) Applying the principle of insufficient reason, we calculate expected values by assigning equal probabilities to each state. If we represent A and B as equally good, then the expected value of B ($\frac{1}{2}k + \frac{1}{2}$) is higher than the expected value of A ($\frac{1}{2}k$). Hence, B is (subjectively) rationally required and A is (subjectively) rationally impermissible. If, on the other hand, we represent A as better than B , then both alternatives have expected values of $\frac{1}{2}$. Hence, A and B are both (subjectively) rationally permissible. Because both ways of representing the values are consistent with the description “ A is at least as good as B ,” each is equally plausible. So, under partition method #2, two equally plausible ways of representing the alternatives’ value yield inconsistent deontic statuses. The same problem arises for partition method #3 (because it includes the state “ B is at least as good as A ”) and partition method #5 (because it includes the state “ A is either better than, worse than, or equally as good as B ”).

Because partition method #1 does not suffer from any of the problems that plague partitions #2 through #5, it is the best partition method. And because there is a single best method of partitioning the states in EI cases for the purposes of applying the principle of insufficient

reason and the MEV norm, the Multiple Partitions Objection fails against the principle of insufficient reason when that principle is applied (together with the MEV norm) in EI cases.

2.3.4 – *The Interval Scale Objection*

The second objection to applying the principle of insufficient reason (together with the MEV norm) to subjective rightness in EI cases is that it requires comparisons in terms of the relevant value to be measurable on an interval scale (Rapoport 1989, 58; Resnik 1987, 36-37). There may, however, be some EI cases in which the relevant value is not interval-measurable. Call this the “Interval Scale Objection.”

Briefly put, *interval scales* provide information about the magnitude of difference between items’ value. On the Fahrenheit temperature scale, for example, the difference between 70°F and 80°F is the same as the difference between 18°F and 28°F. Similarly, the difference between 100°F and 30°F is larger than the difference between 65°F and 70°F. In mathematical terms, interval scales are invariant up to positive linear transformations. In other words, multiplying each measurement by a positive number and adding a constant will preserve the original ranking order and magnitudes. *Ordinal scales*, by contrast, are mere rankings. An example is Moh’s scale of mineral hardness, according to which “*x* is harder than *y* just in case *x* scratches *y* and is not scratched by *y*; and *x* and *y* are equally hard if neither scratches the other or they both scratch one another” (Chang 2002a, 28; see also Griffin 1986, 93). The fact that Moh’s scale ranks diamonds at 10 and quartz at 7 does not allow us to infer anything about *how much* harder diamonds are than quartz.⁴⁵ In mathematical terms, ordinal scales are invariant up to

⁴⁵ To the contrary, Moh’s scale is erratically nonlinear. So, for example, “the progression from calcite to fluorite (from 3 to 4 on the Mohs scale) reflects an increase in hardness of approximately 25 percent,” while “the progression from corundum to diamond, on the other hand (9 to 10 on the Mohs scale), reflects a hardness increase of more than 300 percent.” (Encyclopædia Britannica)

positive monotone transformations. This is a less stringent limitation than for interval scales,⁴⁶ because there is less information to preserve.

Interval scales and their more powerful cousin, *ratio scales*, are cardinal scales. A ratio scale is a cardinal scale because, like an interval scale, it includes information about magnitudes of difference between items. Ratio scales moreover include a non-arbitrary zero point. As a result, ratio scales permit us to make assertions about ratios. This includes both value ratios (e.g., that 20 kilograms is twice as heavy as 10 kilograms) and interval ratios (e.g., that the difference between 20 kilograms and 10 kilograms is twice as much as the difference between 10 kilograms and 5 kilograms). In mathematical terms, ratio scales may be multiplied by a positive constant while preserving the original information about rank orders, magnitudes, and ratios.

Scales should not be confused with *values*. The value of temperature is measurable on a ratio scale (e.g., 311 Kelvin), an interval scale (e.g., 100 degrees Fahrenheit or 38 degrees Celsius), and an ordinal scale (e.g., today was hotter than yesterday). Nor should any ontological significance be attached to the *unit* used in any particular scale. Units in ratio and interval scales are arbitrary—they do not reflect anything about the structure of the value being measured.⁴⁷

Utility of expected utility theory is an interval-measurable value. If I were working within a model of practical reason that took utility as its sole relevant value, then the Interval Scale Objection would be irrelevant. Subjective rightness would never be determined by a value that admitted only of merely ordinal comparisons. But under soft comparativism, a relevant value may be *any* value, and a “value” is liberally defined as anything in terms of which a value relation (better than, worse than, or equally good) may hold (see Ch. 1). This leaves open the

⁴⁶ It includes, for example, logarithmic transformations.

⁴⁷ The *zero point* in a ratio scale, on the other hand, is non-arbitrary. An example: the “distance” between 34 Kelvin and 35 Kelvin is arbitrary; 0 Kelvin is not.

possibility of an EI case where the relevant value—happiness or beauty, perhaps—is not measurable on an interval scale.

There are two possible responses to the Interval Scale Objection. First, we may accept that the principle of insufficient reason and the MEV norm ground alternatives' subjective deontic statuses in EI cases where the relevant value is interval-measurable, and provide a different account of what grounds subjective deontic statuses in EI cases where the relevant value is at best ordinal-measurable.⁴⁸ (Perhaps a different rational norm, together with whatever facts or beliefs that norm takes as inputs.) Second, we may deny that the principle of insufficient reason and MEV norm are operative in *any* EI cases—either those where the relevant value is interval-measurable, or where the relevant value is merely ordinal-measurable.

The strongest argument I can think of for favoring the second response over the first is that it is ad hoc to maintain that one rational norm applies to interval-measurable values and a different rational norm applies to ordinal-measurable values. If there is a single, plausible rational norm that is capable of applying to both interval-measurable and ordinal-measurable values, then principles of parsimony and consistency favor accepting that norm as genuine.

I will argue in the next section that there is such a norm: Maximin. Nevertheless, I do not think that the existence of a single norm that may plausibly govern all EI cases gives us sufficient reason to reject the principle of insufficient reason and the MEV norm as applicable to the subset of EI cases whose relevant value is interval-measurable. The reason for this is that the need for an interval-measurable value comes not from the principle of insufficient reason, but

⁴⁸ I use the term “ordinal-measurable” to mean that an ordinal scale is the strongest scale that may be used to measure the value in question. I use the term “interval-measurable” in the same way. An ordinal-measurable value is thus not measurable with an interval or ratio scale. An interval-measurable value is measurable on either an interval or ordinal scale, but not on a ratio scale.

from the MEV norm. Even in run-of-the-mill decisions with no apparent incommensurability, the MEV norm may ground deontic statuses only if the relevant value is interval-measurable.

Here is an example (inspired by Schotter 2008, 134-135) illustrating why the MEV norm cannot operate in decisions where the relevant value is ordinal-measurable. It takes the form of a *reductio ad absurdum*, first assuming that the MEV norm does operate in decisions where the relevant value is ordinal-measurable, then deriving a contradiction. Suppose that I am meeting a colleague for dinner near my old neighborhood in Madison, and we must decide where to eat. I'd like a nice meal, so as far as I'm concerned, deliciousness is the value relevant to our decision. Let's further assume (for the purposes of the example) that deliciousness is merely ordinal-measurable; it does not admit of interval-measurable comparisons. Here is a ranking, by deliciousness, of several eateries within walking distance of my old duplex. Blue Moon (a dive bar that inexplicably serves food) = 1; Milio's (uninspiring sandwiches) = 2; Brasserie V (good food, Belgian beers on tap) = 3. The particular unit and scale are arbitrary. For an ordinal ranking, any transformation that preserves the original rank order—i.e., any positive monotone transformation—represents the same state of affairs. So the following is also an accurate numerical representation of the eateries' values: Blue Moon = 20; Milio's = 40; Brasserie V = 70. Finally, suppose that for some reason we must choose between: (1) Milio's for sure; and (2) a 50:50 chance of going either to Blue Moon or Brasserie V. According to the first set of values {1,2,3}, the expected deliciousness of Milio's is 2, as is the expected deliciousness of the Blue Moon-Brasserie V gamble ($\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 3 = 2$). Under the MEV norm, either alternative is permissible. But according to the second set of values {20, 40, 70}—which represents the same ordinal ranking—the expected deliciousness of Milio's is 40, while the expected deliciousness of the Blue Moon-Brasserie V gamble is 45 ($\frac{1}{2} \cdot 20 + \frac{1}{2} \cdot 70 = 45$). Now under the MEV norm,

Milio's is impermissible and the gamble is obligatory! The inconsistency may be avoided by instead applying an interval scale, but that is not an option for values that are merely ordinal-measurable. Hence, the MEV norm does not determine alternatives' deontic statuses when the relevant value is merely ordinal-measurable. As Michael Resnik, writing in the context of expected utility theory, points out, this makes intuitive sense:

Ordinal scales represent only the relative standings of the outcomes; they tell us what is ranked first and second, above and below, but no more. In a decision under risk it is often not enough to know that you prefer one outcome to another; you might also need to know whether you prefer an outcome *enough* to take the risks involved in obtaining it. (1987, 81-82)

If the MEV norm grounds alternatives' subjective deontic statuses in run-of-the-mill decisions, then it does so only when the relevant value is interval-measurable. A different rational norm must govern run-of-the-mill decisions when the relevant value is ordinal-measurable. And if two different rational norms govern interval-measurable and ordinal-measurable values in run-of-the-mill decisions, then it is equally plausible that this should be the case for EI cases. There is nothing special about EI cases that makes it uniquely problematic to posit two separate rational norms in that domain. I will not argue here that the MEV norm is operative in run-of-the-mill decisions. I will simply note that the MEV norm is both extremely popular—it enjoys virtually unanimous acceptance by decision theorists, but its popularity extends far beyond that subfield—and intuitively plausible.

The Interval Scale Objection is a response to the proposal that the principle of insufficient reason and the MEV norm jointly ground alternatives' subjective deontic statuses in EI cases. Here are the takeaways from our discussion of it. First, even if the principle of insufficient reason and the MEV norm ground subjective deontic statuses in some EI cases, they cannot do so in all EI cases. They may only do so in EI cases where the relevant value is interval-measurable. So

even if we accept them as operative in EI cases, our work is not done. We still need to find another rational norm that plausibly operates in EI cases where the relevant value is ordinal-measurable. Second, the Interval Scale Objection serves as grounds for rejecting the principle of insufficient reason and the MEV norm in *all* EI cases only to the extent that it serves as grounds for rejecting the MEV norm in other decision-making contexts. Because the MEV norm is so extraordinarily popular, any effort to reject it wholesale must meet a high burden of persuasion. The Interval Scale Objection by itself does not meet that burden. Of course, the MEV norm's popularity does not in itself provide any reason to retain it. But it is so popular because it is highly plausible—it seems to get the right result in a wide variety of decision scenarios, in a variety of domains of practical reason. The fact that it may apply only to comparisons made in terms of interval-measurable values does not itself strike me as sufficient grounds for concluding that it *never* grounds alternatives' subjective deontic statuses.

One final point: although the Interval Scale Objection does not provide sufficient grounds for rejecting the MEV norm, it does slightly weaken the case I initially made for accepting the MEV norm as operative in EI cases. Near the beginning of this section, I made the following claim:

The fact that the norm governing *objective* deontic statuses may be reformulated in MEV terms provides another reason for starting our search for the determinant of *subjective* deontic statuses with the MEV norm. If both objective and subjective deontic statuses may be plausibly grounded in a single norm, then considerations of parsimony and consistency provide an additional reason (beyond its plausibility) to accept that norm as genuine.

The suggestion here is that the MEV norm offers the promise of a single, plausible norm capable of determining all deontic statuses in EI cases—both objective and subjective. The Interval Scale Objection tempers that promise. The MEV norm is capable of determining both objective and subjective deontic statuses, but only in decisions where the relevant value is interval-measurable.

Although it is capable of grounding *all* deontic statuses in the EI cases in which it applies, it is not capable of grounding deontic statuses in *all* EI cases.

2.3.5 – *The Equiprobability Objection*

I turn now to the Equiprobability Objection to the proposal that the principle of insufficient reason and the MEV norm jointly ground subjective deontic statuses in EI cases. While the previous two objections may be at least partially set aside, this objection is decisive. On the basis of the Equiprobability Objection, we should conclude that the principle of insufficient reason is not applicable to EI cases. Because there is no other plausible ancillary assumption that can take the place of the principle of insufficient reason, it follows that the MEV norm is not operative in EI cases. Thus, we must find another rational norm that plausibly governs EI cases.

While the Interval Scale Objection targets the MEV norm, the Equiprobability Objection targets the principle of insufficient reason. As a refresher, here is how I've stated that principle:

principle of insufficient reason: If an agent's beliefs provide no reason to suppose that one state is more probable than another, then the MEV norm treats the states as equally probable.

The principle of insufficient reason is sometimes justified as an instantiation of an older, more general principle of suspending judgment when one's reasons are in equipoise.⁴⁹ This principle of suspending judgment comes from epistemology, and so is typically formulated as a rational constraint on beliefs: when one's reasons equally support two (or more) mutually exclusive propositions, one ought to suspend judgment as to which proposition is true. When formulated in this way, the principle of suspending judgment is a rational constraint on beliefs.

⁴⁹ See, e.g., Dubs 1942, 129 (“[T]his principle [of insufficient reason] means that where ignorance is complete, our expectations for either of two events should ideally be equal, i.e., *we should keep our judgment in entire suspense.*” (emphasis added)). More on other justifications for the principle of insufficient reason below.

As we did earlier with the principle of insufficient reason, it is more useful for my purposes to reformulate it as an ancillary assumption about the grounding of subjective rightness:

principle of suspending judgment: When an agent's beliefs equally support two (or more) mutually exclusive propositions, the grounding of the subjective rightness of her decisions does not include any of those propositions.

In this context, the principle of suspending judgment may be defended as following from the notion of subjective rightness, via a simple *modus tollens* argument:

- (1) If subjective rightness is grounded in a proposition whose truth and falsity are equally supported by the agent's beliefs, then subjective rightness is grounded in information of which the agent is ignorant.
- (2) Subjective rightness is not grounded in information of which the agent is ignorant.
- (3) Therefore, subjective rightness is not grounded in a proposition whose truth and falsity are equally supported by the agent's beliefs.

Subjective rightness is grounded in an agent's beliefs. Subjective rightness, in contrast to objective rightness, is not grounded in information of which the agent is ignorant (premise (2)). If it were, it would not plausibly be able to play the role that it does in assignments of blameworthiness and in guiding agents' decisions. And surely if I have equal reason (or no reason) to believe and disbelieve some proposition, I am ignorant of that proposition (premise (1)). Therefore, since subjective rightness is not grounded in information of which the agent is ignorant, then it is not grounded in a proposition when the agent has equal reason (or no reason) favor its truth or falsity. Subjective rightness effectively suspends judgment as to which proposition is true. Or, if that language feels too metaphorical, we may instead say that subjective rightness is consistent with suspending judgment as to which proposition is true.

In the case of the principle of insufficient reason—which, again, is purportedly an instantiation of the more general principle of suspending judgment—the relevant propositions are: (1) that state 1 is more probable than states 2 through n ; (2) that state 2 is more probable than

either state 1 or states 3 through n ; etc. In the context of a two-alternative EI case, the relevant propositions are:

1. that $[A$ is better than $B]$ is more probable than either $[A$ is worse than $B]$ or $[A$ is equally as good as $B]$
2. that $[A$ is worse than $B]$ is more probable than either $[A$ is better than $B]$ or $[A$ is equally as good as $B]$
3. that $[A$ is equally as good as $B]$ is more probable than either $[A$ is better than $B]$ or $[A$ is worse than $B]$

According to the principle of insufficient reason, the appropriate way to suspend judgment between these propositions is by taking each state to be equally probable. By taking each state as equally probable, the MEV norm avoids any reliance on assumptions about which proposition is true. Subjective rightness, when jointly grounded in the principle of insufficient reason and the MEV norm, is consistent with suspending judgment about which proposition is true.

Propositions #1 through #3 are not logically exhaustive. For a logically exhaustive set of propositions, we would need to add:

4. that $[A$ is better than $B]$, $[A$ is worse than $B]$, and $[A$ is equally as good as $B]$ are equally probable.

But the principle of suspending judgment does not require that the relevant propositions be logically exhaustive; it requires merely that they be mutually exclusive. (This contrasts with constraints on partitioning states for the MEV norm. The states whose probabilities the MEV norm takes as inputs do need to be logically exhaustive, so that their probabilities sum to 1.)

The Equiprobability Objection accepts the principle of suspending judgment, but denies that the principle of insufficient reason is the correct instantiation of that principle to apply to decisions under ignorance (like EI cases). Here is the objection, presented as a logically valid argument:

- (1) In a two-alternative EI case, the agent's beliefs equally support the mutually exclusive propositions #1 through #4.
- (2) If an agent's beliefs equally support two (or more) mutually exclusive propositions, then the grounding of the subjective rightness of her decisions does not include any of those propositions.
- (3) If the principle of insufficient reason partially grounds subjective rightness in a two-alternative EI case, then the ground of subjective rightness includes proposition #4.
- (4) Hence, in a two-alternative EI case, the principle of insufficient reason does not ground the subjective rightness of the agent's decisions.

Premise (2) is the principle of suspending judgment. Premise (3) follows from the proposal, under consideration, that the MEV norm and the principle of insufficient reason jointly ground subjective rightness in EI cases. According to the principle of insufficient reason, the MEV norm in EI cases treats the states as equally probable. In a two-alternative EI case, this means that the MEV norm takes proposition #4 as its input. This, in turn, means that subjective rightness is grounded in (*inter alia*) proposition #4.

The crucial premise is (1). EI cases are decisions under ignorance. Since the alternatives appear incommensurable, the agent's beliefs provide no reason to suppose that one value relation is more likely to hold than another. Put differently, her beliefs do not support the inference that one value relation is more probable than another. But neither do they support the inference that the value relations are equiprobable. She has no more reason to suppose that the value relations are equiprobable than she does to suppose that one value relation is more probable than another.⁵⁰

⁵⁰ See Peterson 2009, 55: "Here is another objection to the principle of insufficient reason: If one has *no* reason to think that one state is more probable than another, it seems strange to conclude that they are *equally* probable. If one has *no* reason to think that some state is *twice* as probable as another, we would not conclude that it is twice as probable as the other. More generally speaking, the problem is that it seems *completely arbitrary* to infer that all states are equally probable. Any other distribution of probabilities seems to be equally justified (that is, not at all)." See also Resnik 1987, 37: "If there is no reason for assigning one set of probabilities rather than another, *there is no justification for assuming that the states are equiprobable either.*" Both authors also raise this issue as a grounds for doubting that the principle of insufficient reason

Taking the value relations as equiprobable requires an ancillary assumption: the principle of insufficient reason. This ancillary assumption is allegedly justified as an instantiation of the principle of suspending judgment. Ironically, that principle itself forms the basis of an argument *against* the principle of insufficient reason—or, more specifically, an argument that the principle of insufficient reason does not apply to EI cases.

The Equiprobability Objection diagnoses the principle of insufficient reason as a case of suspending judgment among the wrong set of propositions. According to the principle of insufficient reason, subjective rightness suspends judgment between the propositions that one state is more probable than another. Taking each state as equally probable is a way of doing so. According to premise (1) of the Equiprobability Objection, subjective rightness instead suspends judgment between *probability distributions*—that is, between a mutually exclusive and exhaustive set of propositions that, by dint of being logically exhaustive, represent every possible distribution of probabilities between the states. Again, the motivation for including proposition #4 is not that its inclusion makes the set of propositions logically exhaustive—the principle of suspending judgment does not require a logically exhaustive set of propositions. Instead, the contention is that subjective rightness suspends judgment with respect to proposition #4 for exactly the same reason that it suspends judgment with respect to propositions #1, #2, and #3. It is equally as supported by the agent's beliefs as the other, mutually exclusive propositions #1 through #3.

Perhaps, one might reply, the principle of insufficient reason may be saved by justifying it on some other grounds—that is, not on the grounds that it represents a form of suspending

applies to decisions under ignorance. My thinking about the Equiprobability Objection owes much to both discussions.

judgment.⁵¹ Indeed, there are at least two other justifications that have been offered for the principle of insufficient reason. One is empirical confirmation: “For example, consider that you are throwing a die. Since you do not believe that one side is more likely to occur than any other, you regard all probabilities as equal. And indeed, repetitive throwing shows that this is correct” (Sinn 1980, 493). The other is axiomatic, and involves arguing that the principle of insufficient reason follows from certain axioms commonly accepted in expected utility theory—e.g., transitivity and independence (see Sinn 1980).

Empirical confirmation, at least, is not available for EI cases. Because EI cases involve normative ignorance, it is not possible to empirically verify whether the principle of insufficient reason assigned the correct probability to each state. The alternatives will appear just as incommensurable to the agent after she makes her choice as they did before. There is no die for her to look at to confirm or disconfirm the principle’s accuracy.

The better response, however, is that the Equiprobability Objection stands even if there are other justifications for the principle of insufficient reason. Whatever other arguments we may find supporting the principle of insufficient reason, it will still be the case that the beliefs of an agent facing an EI case no more support an inference that each value relation is equiprobable than they support an inference that any particular value relation is more probable than another. If the principle of suspending judgment is, as I have argued, a genuine constraint on the grounds of subjective rightness, then the principle of insufficient reason does not ground subjective rightness in EI cases.

⁵¹ Reznik 1987, 37 suggests that this may be possible: “If every probability assignment is groundless, the only rational alternative is to assign none at all. *Unless some other rationale can be provided for the principle of insufficient reason*, that means that we should turn to some other rule for making decisions under ignorance.” (emphasis added)

To accept the principle of insufficient reason as a partial ground of subjective rightness in EI cases, we would need to either reject premise (1) or premise (2) of the Equiprobability Objection. That is, we would need to either: (1) reject the notion that the value relations' equiprobability is as unsupported by the agent's beliefs as the assignment of a heightened probability to any particular value relation; or (2) reject the principle of suspending judgment.

It is hard to see why we would reject premise (1). There are certainly cases where the agent is uncertain as to which value relation holds between her alternatives, but in which she is leaning towards one or another. (I might be uncertain whether a donation to Oxfam or Doctors Without Borders will maximize moral value, but have grounds for thinking it more likely that the Oxfam donation will.) But EI cases just are those choices in which this is not the case. The agent's uncertainty as to which alternative is most valuable is so total that they appear incommensurable to her. It is also hard to see why we would reject premise (2), since it follows so closely from the notion of subjective rightness.

There is one context in which the principle of insufficient reason *is* a plausible instantiation of the principle of suspending judgment. This is if we take on the additional assumption that determinations of subjective rightness *necessarily* require one probability or another. If that is so, then subjective rightness *cannot* suspend judgment among probability distributions. The best that it can do is suspend judgment as to which value relation (that is, which state) is most probable by taking them to be equally probable.

This assumption is built into some models of practical rationality. For example, Bayesianism's "precise probabilism" holds that "no evidence is ever silent on any proposition;" hence, "[o]ne can never suspend judgment" on the question of which probability distribution holds between a set of mutually exclusive and logically exhaustive states (Bradley 2016, citing

Huber 2009). That assumption may make good sense in the context of an ideal agency theory of epistemology.⁵² And indeed, it is widely accepted in Bayesian epistemology, where it is more often called the “principle of indifference.” The assumption seems less well-motivated, however, as part of a general and less restrictive theory of how subjective rightness is determined for actual agents.

We would also have reason to seriously entertain this assumption if there were no other rational norm, aside from the MEV norm, that could plausibly determine subjective rightness in EI cases. The MEV norm takes states’ probabilities as inputs. So, if subjective rightness is determined by the MEV norm, then subjective rightness cannot suspend judgment as to those states’ probabilities.

Fortunately, as I will argue in the next section, there *is* another rational norm—Maximin—that plausibly grounds subjective rightness in EI cases. John Rawls, when making his case that the maximin norm applies to decisions made behind the veil of ignorance, suggested that in decisions under ignorance (which include both decisions under Rawls’ veil of ignorance⁵³ and EI cases), “it is unreasonable not to be skeptical of probabilistic calculations unless there is no other way out” (Rawls 1971, § 26). The Equiprobability Objection provides an argument as to why this is so. Fortunately, for EI cases, as for decisions behind the veil of ignorance, there is another way out—another rational norm that plausibly grounds subjective rightness. Because of the problems that the Equiprobability Objection poses for the principle of insufficient reason, and because there is another rational norm that can plausibly ground alternatives’ subjective deontic statuses in EI cases *without* relying on the principle of insufficient reason, we should reject the

⁵² Then again, it may not. Bradley 2016 suggests that the assumption of precise probabilism is itself “a nice way of seeing the strangeness of the precise probabilist’s attitude to evidence.”

⁵³ Behind the veil of ignorance, “a knowledge of likelihoods is impossible, or at best extremely insecure” (Rawls 1971, §26).

proposal that the principle of insufficient reason and MEV norm jointly ground alternatives' subjective deontic statuses in EI cases.

2.4 – Maximin

2.4.1 – *The MAPV Norm, Maximin, Maximax, and Medianmax*

The hunt for a rational norm that plausibly grounds alternatives' subjective deontic statuses in EI cases continues. I started this hunt with the premise that the most obvious norm is one that maximizes value. This claim is itself hard to justify, and I will not go to any great lengths to do so here. As I noted in Chapter 1, there are all kinds of choice scenarios in which it plausibly is rationally permissible—perhaps even obligatory—for an agent to fail to maximize (or optimize) value. But the mechanisms that ground alternatives' deontic statuses in those scenarios (e.g., exclusionary reasons) can, by and large, be relegated to the role of deontic constraints.⁵⁴ It is something close to a conceptual truth that, once deontic constraints have been accounted for, rational action just is action that maximizes value (or perhaps optimizes value in OI cases, where the notions of maximization and optimization come apart).⁵⁵

⁵⁴ Even supererogation and Michael Slote's version of satisficing (see above, n.25) can, I think, be construed as a kind of deontic constraint. Where a non-value-maximizing alternative is permissible, it is because there is some normative fact, or some feature of normativity, that suspends practical reason's default of maximization and grounds alternatives' deontic statuses in some other way. I say "default," but the claim I have just made is consistent with supererogation and/or satisficing being permissible in the majority of choices. If one has a robust theory of supererogation and/or satisficing such that those norms are operative in most choices, then the "default" of maximizing turns out to be operative only in a minority of cases; it is the catch-all norm that is operative when other norms are not. What I have to say on the topic of maximizing norms will thus hopefully still be of some interest to proponents of these norms.

⁵⁵ The notable exception is parity, for philosophers who accept that value relation as genuine. Parity's proponents argue that either of two on-a-par alternatives is permissible even though neither maximizes value (because neither is better than the other). But this exception is inapplicable to EI cases, since parity entails incommensurability. Moreover, for philosophers who take it as axiomatic that rational action (absent deontic constraints) just is action that maximizes value, this is simply a reason to reject parity as a genuine value relation.

An alternative's actual value is the value that it in fact has. An alternative's expected value is a probability-weighted average of its possible values. Either of these is maximizable. When an agent is ignorant of her alternatives' actual value, it is sometimes still possible to maximize their expected value. I have argued that neither a norm of maximizing actual value (the MAV norm) nor a norm of maximizing expected value (the MEV norm) grounds alternatives' deontic statuses in EI cases. Fortunately, there remain other ways to make meaningful claims about maximizing value. When the relevant probabilities do not exist, as is true of EI cases, it may still be possible to maximize the *non-weighted* average of an alternative's possible values. (Call the rational norm that prescribes this the "MAPV norm," for maximizing *average possible value*.) Or we may dispense with averaging altogether, and maximize alternatives' *minimum possible value*. (Call the norm that prescribes this "Maximin.") Alternatively, we might maximize alternatives' *maximum possible value* (as prescribed by "Maximax"). Finally, we might consider a "Medianmax" norm, according to which an agent ought maximize alternatives' *median possible value*. In the wake of the MEV norm's failure, then, we have at least four other value-maximizing norms that might ground subjective deontic statuses in EI cases: the MAPV norm, Maximin, Maximax, and Medianmax.

Here is the methodology that I will follow for the rest of the chapter. I will assume that there is, at the least, a strong presumption that absent deontic constraints, rational action just is action that maximizes the value relevant to the choice. In other words, there is a strong presumption that the rational norms that ground alternatives' objective and subjective deontic statuses hold that an alternative is rationally permissible if and only if it maximizes value, relative to the agent's other alternatives. Having ruled out the MAV norm and MEV norm, the most straightforward value-maximizing norms that might ground subjective deontic statuses in

EI cases are the four candidate norms from the preceding paragraph. In what follows, I will examine whether a compelling case can be made that one of these norms, rather than the others, grounds subjective deontic statuses in EI cases. If a compelling case cannot be made for any of the candidate norms, then we will be forced to either (1) search for a more exotic form of value-maximizing norm,⁵⁶ or (2) rebut the strong presumption and conclude that subjective deontic statuses in EI cases are not grounded in a value-maximizing norm. I will argue that Maximin is the operative norm, and so we will not be forced into this uncomfortable dilemma.

Both the MAPV norm and Medianmax require, like the MEV norm, that the relevant value be interval-measurable.⁵⁷ So even if we determined that MAPV norm or Medianmax is operative in some EI cases, we would still need to press on to find a plausible, value-maximizing norm that is capable of grounding subjective deontic statuses when the relevant value is merely ordinal-measurable. Maximin and Maximax, on the other hand, require merely ordinal inputs.

⁵⁶ For example, we might consider norms that maximize a value other than the value relevant to the choice. One such norm is the Minimax Regret norm (see Savage 1954), according to which the agent ought choose the alternative that minimizes the agent's *maximum possible regret* at having failed to choose the alternative that maximizes actual value. The Minimax Regret norm is a value-maximizing norm, but it maximizes a different value than what I have been calling the "relevant value." It instead maximizes *lack of regret*.

⁵⁷ This can be shown by replacing the MEV norm with the MAPV norm and Medianmax in the example from section 2.3.4. The choice, recall, is between (1) Milio's for sure; and (2) a 50:50 chance of going either to Blue Moon or Brasserie V.

First, suppose that the MAPV norm is operative. According to the first set of values {1,2,3}, the average possible deliciousness of Milio's is 2, as is the average deliciousness of the Blue Moon-Brasserie V gamble ($(1+3)/2 = 2$). Under the MAPV norm, either alternative is permissible. According to the second set of values {20, 40, 70}, the average possible deliciousness of Milio's is 40, while the average possible deliciousness of the Blue Moon-Brasserie V gamble is 45 ($(20 + 70)/2 = 45$). Under the MAPV norm, the gamble is now obligatory. As with the MEV norm, then, using an ordinal scale to measure the relevant value results in the MAPV norm rendering inconsistent deontic statuses.

Second, suppose that Medianmax is operative. Assuming that we take the median of an even set of values in the usual way by averaging the two middle values, the median possible values in this choice wind up being identical to the average possible values. Accordingly, Medianmax, too, requires that the relevant value be interval-measurable.

They are thus capable of grounding subjective deontic statuses in all EI choices, regardless of whether the relevant value is measurable on a cardinal scale.

It turns out that in a two-alternative EI case, all four norms render the same set of deontic statuses: either alternative is subjectively rationally permissible. Here is a walkthrough of how this works, assigning values as we did earlier (in Table 2.1):

Table 2.7

	A is better than B	A is worse than B	A is equally as good as B
A	2	1	k
B	1	2	k

With respect to the MAPV norm, either alternative is (subjectively rationally) permissible because both alternatives have the same average possible value: $(3 + k)/3$. With respect to Maximin, Maximax, and Medianmax, either alternative is (subjectively rationally) permissible because each alternative has the same set of possible values: $\{1, 2, k\}$. Accordingly, A and B have identical medians, regardless of how we fix the value of k . Similarly, A and B have identical minimum possible values (either 1 or k , if k is lower than 1) and identical maximum possible values (either 2 or k , if k is higher than 2⁵⁸) regardless of how we fix the value of k .

It is also possible to describe how Maximin and Maximax would operate in EI cases without resorting to any numerical representation of the alternatives' values. The agent's ignorance in two-alternative EI cases, as reflected by the way in which I've partitioned the states, pertains to which value relation holds between A and B . Equivalently, the agent is ignorant as to A and B 's *relative values*—their value relative to one another. Maximin and Maximax merely require ordinal inputs. Hence, for the purposes of these rational norms, we may say that the agent is ignorant as to A and B 's ordinal ranking with respect to the relevant value. She is, in other

⁵⁸ Assigning k 's value as either lower than 1 or higher than 2 seems nonsensical to me. But my point is that the norms yield the deontic statuses they do *regardless of* what value we assign to k .

words, ignorant as to whether the relevant value ranks A above B , ranks B above A , or ranks A and B together. For either alternative, there is a possible state of the world in which it is ranked lowest (i.e., is the worse alternative), a state in which it is ranked highest (i.e., is the better alternative), and a state in which, by dint of being equally as valuable as the other alternative, it is simultaneously in the lowest and highest ordinal rank. Under Maximin, both alternatives are (subjectively rationally) permissible because their minimum possible values are identical. For each alternative there is a state in which it has the lowest ordinal ranking of any available alternative. Under Maximax, both alternatives are (subjectively rationally) permissible because their maximum possible values are identical. For each alternative there is a state in which it has the highest ordinal ranking of any available alternative. So, we can describe how Maximin and Maximax operate in terms of ordinal ranking, without any numerical representation. This is not, however, possible for the MAPV norm, since that norm grounds subjective deontic statuses in an arithmetic mean.

The four norms will also yield identical deontic statuses in EI cases with more than two alternatives where, as in the two-alternative case above, there is *complete* apparent incommensurability between alternatives. By “complete” apparent incommensurability, I mean that for any alternative, A , A appears incommensurable with every other alternative. If apparent incommensurability is complete, no alternative appears commensurable with any other alternative. So long as an EI case involves complete apparent incommensurability, then according to any of the four candidate norms, any alternative is subjectively rationally permissible. But it would be premature of us to conclude that, for the purposes of ascertaining alternatives’ subjective deontic statuses, we need not bother with determining which of the four candidate norms is operative in EI cases. For in cases of *partial* apparent incommensurability,

where the agent knows some but not all of the value relations that hold between alternatives, the norms sometimes render different prescriptions.

Consider, for example, a three-alternative $\{A, B, C\}$ EI case in which: (1) A appears incommensurable with either B or C , but (2) the agent perceives B as slightly better than C . Perhaps the National Book Foundation's panelists are deciding, on the basis of literary merit, which newly published novel should be awarded the National Book Award. The Foundation's panelists have narrowed the list to two historical fiction novels (B and C) and one magical realist novel (A). Pressing themselves further, they are able to judge B as slightly better than C . But despite their best efforts, they cannot confidently judge whether A is better or worse than B or C . A is so qualitatively dissimilar—it manifests the value of literary merit in such different ways than B or C —that it appears wholly incommensurable with them. (Assume for the purposes of the example that this is a case of EI rather than OI.)

Under Maximin, the judges ought to give the award to B . B maximizes minimum possible value, because under any possible state of the world, there is another alternative— C —that is worse. While there are possible states of the world in which A and C are each ordinally ranked lowest, there is no possible state of the world in which B is ordinally ranked lowest. Hence, B maximizes minimum possible value. Under Maximax, on the other hand, the judges may give the award to either A or B . A may turn out to be the most valuable alternative, or B may turn out to be the most valuable alternative, but in no possible state of the world is C the most valuable alternative. There is always another alternative— B —that is better than C . Accordingly, only A or B maximizes possible value (on an ordinal scale). Because the norms sometimes, as in this case, yield inconsistent subjective deontic statuses, we cannot remain agnostic as to which grounds

subjective deontic statuses in EI cases. We must decide which of the four candidate norms is operative in EI cases.

2.4.2 – *Why Maximin is the Operative Norm in EI Cases*

Of the four candidate norms, Maximin and Maximax are the only two that I have seen discussed in print. Maximin is the more popular of the two and, perhaps because of this, Maximax is sometimes dismissed with little to no argument.⁵⁹ I, too, find Maximin to be the most intuitively plausible of the four norms. But because not everyone agrees,⁶⁰ it is incumbent on Maximin’s proponents to defend it against the other candidate norms.

The most well-known defense of Maximin as a ground of subjective deontic statuses is in Rawls 1971.⁶¹ But Rawls’ argument for Maximin depends on features peculiar to the Rawlsian original position. Rawls argues that there are “three chief features of situations that give plausibility to” the “unusual” rational norm of Maximin (1971, § 26). “The paradigm situation for following the maximin rule,” Rawls writes, “is when all three features are realized to the highest degree.” Maximin “does not ... generally apply, nor of course is it self-evident.” The first feature, which I have argued *is* present in EI cases (see above, § 2.3.1), is that the agent’s choice is a decision under ignorance, rather than a decision under risk.⁶² This feature gives us reason to

⁵⁹ Rapoport 1989, 57, for example, brusquely dismisses Maximax as “reflecting ... wishful thinking.”

⁶⁰ See, e.g., Peterson 2009, 47 (“[A]s far as I can tell, [Maximax] is about as reasonable as the maximin rule.”).

⁶¹ Because Rawls analyzes the rationality of choices made behind a veil of ignorance (see 1971, § 24), he is best characterized as discussing subjective, rather than objective, deontic statuses.

⁶² In Rawls’ words: “[S]ince [Maximin] takes no account of the likelihoods of the possible circumstances, there must be some reason for sharply discounting estimates of these probabilities. Offhand, the most natural rule of choice would seem to be [the MEV norm]. Thus it must be, for example, that the situation is one in which a knowledge of likelihoods is impossible, or at best extremely insecure. In this case it is unreasonable not to be skeptical of probabilistic calculations unless there is no other way out, particularly if the decision is a fundamental one that needs to be justified to others.” (1971, § 26)

favor the four candidate norms over the MEV norm, because (as discussed above) the MEV norm is plausible only for decisions under risk.

Here are Rawls' other two features, which he presumably takes to justify the conclusion that Maximin, rather than some other non-MEV candidate norm, determines alternatives' subjective deontic statuses in choices made from his original position:

Second feature: “[T]he person choosing has a conception of the good such that he cares very little, if anything, for what he might gain above the minimum [value] that he can, in fact, be sure of by following the maximin rule. It is not worthwhile for him to take a chance for the sake of a further advantage, especially when it may turn out that he loses much that is important to him.”

Third feature: “[T]he rejected alternatives have outcomes that one can hardly accept. The situation involves grave risks.”

Taken together, these two features describe a choice scenario in which: (1) there is a minimum threshold of value that it is rationally obligatory for the agent to obtain in her choice (an agent can “hardly accept” an alternative below the threshold); and (2) the alternative(s) that is (or are) permissible according to Maximin has (or have) a minimum possible value that meet(s) or exceed(s) that threshold level. It is (subjectively rationally) impermissible for agents to choose an alternative that is not permitted by Maximin, even if it might result in their obtaining greater value than can be obtained under any state with the Maximin-prescribed alternative(s), because, in forgoing the Maximin-prescribed alternative(s), they take on the risk of obtaining an alternative that falls below the minimum value threshold.

Rawls argues that his original position is just this sort of choice scenario. Agents in Rawls' original position must choose between competing principles of justice. According to Rawls, the only alternative that Maximin renders (subjectively rationally) permissible is his lexically ordered two principles of justice. Rawls' two principles together “guarantee ... a satisfactory minimum” (1971, 135). In other words: for each alternative that Maximin renders

(subjectively rationally) permissible, there is no possible state of the world in which that alternative's value falls below the minimum threshold.⁶³ On the other hand, other possible principles of justice (i.e., other alternatives available to the agent) “may lead to institutions [e.g., slavery] that the parties would find intolerable.” In other words: for each alternative that Maximin renders (subjectively rationally) impermissible—e.g., the principle of (average or aggregate) utility—there is at least one state in which that alternative's value falls below the minimum threshold.

Thus, according to Rawls, Maximin (1) correctly categorizes alternatives that meet the minimum value threshold under every possible state as permissible, and (2) correctly categorizes alternatives whose value falls below the minimum threshold under at least one state as impermissible. There are a couple of ways to read Rawls here. The weak reading is that Rawls is making a claim only about Maximin's plausibility: that it plausibly grounds deontic statuses for choices in the original position because the (subjective) deontic statuses that it yields correctly reflect the minimum value threshold that is obligatory in those choices. The strong reading is that Rawls is moreover making a claim about *other candidate norms*' implausibility. For all other candidate norms, on this stronger reading, there is at least one alternative—one principle of justice—that the candidate norm renders (subjectively rationally) permissible even though there is at least one state in which that alternative's value falls below the minimum value threshold, rendering the alternative *ipso facto* (subjectively rationally) impermissible. That is, all other

⁶³ Rawls' argument that Maximin is the operative rational norm applies only to choices in which there is an alternative whose minimum possible value meets or exceeds the minimum value threshold (what Rawls calls the “satisfactory minimum”). If there were no such alternative—i.e., if every alternative's minimum possible value fell below the minimum value threshold—then the minimum value threshold would not be obligatory. It would be merely desirable. But “having the ready alternative of the two principles of justice which” *do* “secure a satisfactory minimum, it seems unwise, if not irrational, for [agents in the original position] to take a chance that these conditions are not realized” (1971, 135).

candidate norms yield inconsistent deontic statuses by rendering at least one (subjectively) impermissible alternative (subjectively) permissible. Under the strong reading of Rawls, we are justified in concluding that Maximin grounds alternatives' subjective deontic statuses in choice scenarios with his "three features" (e.g., the Rawlsian original position) because in those scenarios, Maximin is the only candidate norm that fails to yield inconsistent deontic statuses.

Rawls' second and third features are peculiar to the Rawlsian original position, and it is doubtful that they are present in all (or even most) EI cases. Arguably, the presence of a rationally obligatory minimum threshold of value is best interpreted as a deontic constraint. If this is so, then Rawls' second and third features will not be present in *any* of the cases of apparent incommensurability (EI or OI) that I discuss in this dissertation, since I've set aside choice scenarios with deontic constraints. So even if Rawls' second and third features justify the conclusion that Maximin is the operative rational norm for *choices made in his original position*, they arguably do not justify that conclusion for *EI cases*.

Fortunately there is another way to justify the conclusion that Maximin, rather than one of the other three candidate norms, grounds subjective deontic statuses in EI cases. The alternative justification is based on the fact that Maximin, in contrast to the other three candidate norms, is what we might call is a "sure thing" norm. An agent who chooses according to the prescriptions rendered by Maximin is guaranteed to obtain the value that it maximizes, regardless of which possible state of the world obtains. She will obtain that value "for sure" (or, as decision theorists sometimes say, "come what may"). This is what Rawls means when he states in his second feature that Maximin has a minimum that the agent can "be sure of by following the maximin rule."

Here is a decision matrix to illustrate how this works:

Table 2.8

	state 1	state 2	state 3
<i>A</i>	0	10	20
<i>B</i>	1	14	15
<i>C</i>	8	10	12

In Table 2.8, the alternative that maximizes minimum possible value is *C*. *C* has a minimum possible value of 8, while *A* and *B* have minimum possible values of 0 and 1, respectively.

According to Maximin, *C* is therefore the only (subjectively rationally) permissible alternative.

Moreover, by choosing *C*, the agent is guaranteed to obtain a value of 8, regardless of which state obtains. She might obtain more, but she cannot obtain less.

If an agent chooses according to any of the other three candidate norms, she gives up this sure thing. She may wind up obtaining more value than if she had chosen according to Maximin or, depending on which state obtains, she may wind up with less. Under Maximax, she is (subjectively rationally) obligated to choose *A*, since *A* has the highest maximum value: 20. If either state 2 or state 3 obtains, she will obtain more value (either 10 or 20) than if she had chosen according to Maximin (8); however, if state 1 obtains, she will obtain less value (0) than if she had chosen according to Maximin. Under Maximedial, she is (subjectively rationally) obligated to choose *B*, because its median value (14) is higher than that of the other alternatives (10). Once again, she may wind up with more value (either 14 or 15) than if she'd chosen under Maximin (8), or she may wind up with less (1). Under the MAPV norm, she is (subjectively rationally) permitted to choose any alternative, since every alternative has an average possible value of 10. As before, if she chooses either *A* or *B*, she may wind up with either more or less value than the sure thing guaranteed by Maximin.

If an agent chooses according to Maximin, she obtains a guaranteed value. If she instead chooses according to Maximax, Maximedial, or the MAPV norm, she takes on risk: she may

obtain more value, or she may obtain less. Because it is possible (though not certain) for an agent to obtain more value by choosing according to a norm other than Maximin, it is sometimes said that Maximin is a more *conservative* norm than the others, and that we are justified in concluding that it grounds alternatives' (subjective) deontic statuses only if the choice in question has special features that make a conservative norm plausible. Other times (the thinking goes)—as perhaps when there is an alternative with a maximum possible value that greatly exceeds the value guaranteed by Maximin—we may be justified instead in concluding that a less conservative norm grounds alternatives' (subjective) deontic statuses. Rawls' argument for Maximin is consistent with this line of thinking. He argues that his original position has special features in virtue of which Maximin plausibly grounds alternatives' (subjective) deontic statuses.⁶⁴

I believe that this line of thinking is mistaken. Let us say that an agent *takes on risk* just in case she gives up a sure thing in exchange for the possibility that she will obtain less value. I take it to be an uncontroversial tenet of practical rationality that taking on risk is rational only if doing so is justified: risk (unless unavoidable) requires justification.⁶⁵ There is, in other words, a

⁶⁴ It is because his original position manifest his three features “to a very high degree,” writes Rawls, that “it is rational for the parties to adopt the conservative attitude expressed by [Maximin]” (1971, 133).

⁶⁵ My committee tells me that this tenet is more controversial than I have suggested. Here, for example, is a counterexample offered to me by Peter Vranas. Suppose that an agent must choose between two gambles. Gamble #1 offers outcomes of either \$0 or \$1, with unknown probabilities. Gamble #2 offers outcomes of either \$-1 or \$2, with unknown probabilities. Given that I stand to lose only a trivial amount of value if I take on Gamble #2, it is rationally permissible (the thinking goes) for me to choose Gamble #2. But my own inclination here is to push back. Why should I take on Gamble #2's possibility of losing \$1—why is it rational for me to do so—if I have *no sense at all* of the probability of the potential \$2 gain? If I knew (or could estimate) the probabilities, taking on the possibility of loss could be justified via the MEV norm. (More on this below.) Or if Gamble #2's possibility of trivial loss were paired with a possibility of non-trivial gain, you might think that taking on the possibility of loss is justifiable via Maximax. (More on this below, too.) But why should I take on the possibility of even a trivial loss for the possibility of a trivial gain? My own intuitions, at any rate, lean strongly towards the tenet that doing so requires justification. And my arguments below show why I do not believe

burden of proof on those who would argue that, for some instance of taking on risk, it is (subjectively rationally) permissible for the agent to take on the risk. It follows that Maximin is the “default” norm, because in choosing according to any other rational norm an agent takes on risk. The previous paragraph’s line of thinking wrongly suggests that there is a burden of proof on Maximin’s proponents to justify why the agent ought forego the possibility of obtaining more value (than is obtainable if the agent chooses according to Maximin). But the burden of proof is instead on those who claim that for some choice, a norm other than Maximin grounds alternatives’ subjective deontic statuses. All other norms permit the agent to take on risk; proponents of those norms thus have the burden of proving that taking on the risk is justified.

In so-called decisions under risk, the orthodox view is that taking on risk is justified if and only if the expected value of the risk is at least as great as the value that the sure thing instantiates. In other words, taking on risk is justified if and only if, in a choice between the risk and the sure thing, taking on the risk is (subjectively rationally) permissible under the MEV norm. So, for example, wagering \$1 (i.e., giving up \$1 “for sure”) on the outcome of a horse race is justified if and only if the expected value of the wager is at least \$1 (assuming for the purposes of the example that the relevant value is money and not something that admits of diminishing marginal value).

The MEV norm, in turn, is typically justified—to the extent that appeals to bare intuition and the plausibility of its prescriptions are deemed insufficient—via either long-run arguments or representation theorems. Stated crudely, long-run arguments justify taking on those risks that are permissible under the MEV norm on the grounds that if the agent did so repeatedly she would, in the long run, obtain a value that becomes arbitrarily close to an expected value that is *at least as*

that such a justification—i.e., a justification for taking on the possibility of loss in a decision under ignorance—is possible.

good as the sure thing. Representation theorems assume a set of axioms that are taken to be plausible constraints on rational preferences,⁶⁶ and then show that an agent whose preferences obey the axioms (and who chooses with the aim of satisfying her preferences) acts in conformity with the MEV norm. Representation theorems thus justify taking on those risks that are permissible under the MEV norm on the grounds that an agent who fails to do so necessarily violates one or more of the axioms.

In EI cases and other decisions under ignorance, risk cannot be justified in terms of expected value, because there are no subjective probabilities. Nor can risk be justified on the basis of existing long-run arguments or existing representation theorems. Long-run arguments purport to show that if an agent acts in accordance with the MEV norm, she will obtain a value that converges, in the limit, with the alternative's expected value. But there is no guarantee that an agent who acts in accordance with the MAPV norm, Maximax, or Medianmax will thereby obey the MEV norm. To the contrary, these norms sometimes render different prescriptions than the MEV norm.⁶⁷ Accordingly, the value obtained by an agent who follows these norms does not converge, in the limit, with expected value. It may be possible to construct a novel long-run argument to show that an agent who follows one of these other norms obtains, in the long run, a value that is at least as good as that guaranteed by the sure thing. But the burden of proof is on those norms' proponents to formulate such an argument.

Similarly, representation norms justify all and only those risks that are permissible under the MEV norm. But that set of risks is not identical to the set of risks that are permissible under

⁶⁶ In other words, an agent's set of preferences is rationally permissible if and only if it obeys the axioms.

⁶⁷ This occurs, for example, in the Table 2.8 decision, if we take each state to have a probability of 1/3. Because each alternative has an expected value of 10, any alternative is subjectively rationally permissible under the MEV norm. This is inconsistent with the subjective deontic statuses rendered by Maximin, Maximax, or Medianmax.

the MAPV norm, Maximax, or Medianmax. It might, of course, be possible to construct a representation norm for each of these other norms by adopting a different set of rational preference axioms.⁶⁸ But the burden of proof is on those norms' proponents to both formulate the axioms and defend those axioms as plausible.

I suspect that the unavailability of existing long-run arguments in particular to justify rational norms for decisions under ignorance is the reason that no one ever proposes the MAPV norm. Setting aside the well-known criticisms of long-run arguments,⁶⁹ an alternative's *probability-weighted* average value is plausibly relevant to its (subjective rational) permissibility in part because it represents a value that the agent is virtually assured of obtaining over the long run. A probability-weighted average represents, if not a sure thing, at least a value that is arbitrarily close to an expected value that is at least as good as the sure thing. An alternative's *non-probability-weighted* average value represents no such thing. It is neither a guaranteed value, nor a value that is virtually guaranteed in the long run; it is merely an average of the alternative's possible values. But without the long-run argument or some other justification, there is no obvious reason why that value has any bearing on the alternative's subjective deontic status.

The same can be said of Medianmax, which is also never proposed, and of Maximax, which *is* often mentioned as a candidate norm for the ground of subjective rightness in decisions under ignorance. Because Maximin guarantees a sure thing, the risk that the other candidate norms require agents to take on requires justification. But the circumstances and arguments that we typically take to justify risk are unavailable in EI cases and other decisions under ignorance.

⁶⁸ For an example of this strategy in a different context, see Buchak 2013, Ch. 3. Buchak adopts a different set of axioms than those typically used in expected utility representation theorems, then uses those axioms to construct a representation theorem for her risk-weighted expected utility theory.

⁶⁹ For an overview of these, see Briggs 2017, § 2.1.

Absent some other compelling justification, we should conclude that Maximin is the operative norm for subjective rightness in EI cases.

The best contender for an alternative justification for taking on risk in decisions under ignorance relies on a family of alleged counterexamples to Maximin. These are choices in which there is an alternative with a maximum possible value that greatly exceeds the value guaranteed by Maximin. One particularly clever version combines an arbitrarily high maximum possible value with an arbitrarily small loss (if the agent takes on risk and forgoes Maximin's guaranteed value):

Table 2.9⁷⁰

	state 1	state 2
<i>A</i>	0	<i>n</i>
<i>B</i>	1/ <i>n</i>	1

Surely, the argument goes, there is some positive integer, *n*—make it as high as you like—such that it is at least (subjectively rationally) permissible, and perhaps even (subjectively rationally) obligatory, to choose *A* rather than *B*. That is, there is some point at which the maximum possible value is sufficiently higher than the sure thing guaranteed by Maximin, and the possible loss compared to Maximin's sure thing sufficiently negligible, that it is at least permissible and perhaps obligatory to take on risk.

As a first step in evaluating this argument, we should clarify exactly what it means for something to be a counterexample to Maximin. First, the argument might be interpreted as a refutation of the proposal that Maximin grounds subjective deontic statuses in *all* decisions under ignorance. If sound, the argument negates that proposal by offering a valid counterexample: a

⁷⁰ I owe this particular formulation to Rawls 1971, 136, though it is not original to him. Unfortunately, Rawls does not actually address this counterexample in the cited passage. He simply mentions that it is analogous to a similar worry that might be raised against his difference principle.

decision under ignorance in which Maximin does not ground subjective deontic statuses. (Call the argument, under this interpretation, the “*n*-Value Objection.”)

Second, the argument might be offered in partial support of the stronger conclusion that Maximin does not ground *any* deontic statuses in decisions under ignorance—i.e., that there is no decision under ignorance such that Maximin grounds the alternatives’ deontic statuses. But this conclusion does not follow without adopting further premises. One premise that might support this stronger conclusion is that considerations of parsimony and consistency favor theorizing that a single rational norm is operative for all decisions under ignorance. If combined with the further premise that there is a single rational norm that is capable of plausibly grounding deontic statuses in all decisions under ignorance, the counterexample demonstrates why theoretical considerations favor accepting that other norm (rather than Maximin) as genuine. (Call this expanded version of the argument the “Expanded *n*-Value Objection.”)

The Expanded *n*-Value Objection is unsound, at least when formulated in this way, because both of the additional premises are false. As I have described things, Maximin is a “default” norm that grounds subjective rightness in *all* choices, absent some justification for taking on risk. This entails that if risk is justified for some but not all members of a particular subset of choices (e.g., decisions under ignorance), multiple norms will ground the members of that set of choices. If taking on risk is justified in some but not all decisions under ignorance, for example, Maximin grounds the subjective deontic statuses in all and only those decisions under ignorance in which taking on risk is not justified. Some other norm grounds the subjective deontic statuses when taking on risk *is* justified—and hence permissible—because taking on risk is impermissible under Maximin. Positing a single norm as operative in all choices (absent special circumstances) strikes me as a better-motivated application of parsimony and consistency

considerations than merely preferring that all decisions under ignorance have their subjective deontic statuses grounded in a single norm. If so, then the first additional premise in the Expanded n -Value Objection is false. Because considerations of parsimony and consistency support positing Maximin as a single norm operative in all choices (absent justification for taking on risk), they do not favor theorizing that a single rational norm is operative for all decisions under ignorance.

The second additional premise is also false. The n -Value Objection allegedly shows that there is some value, n , such that Maximin does not plausibly ground the alternatives' subjective deontic statuses in Table 2.9. For the Expanded n -Value Objection to run, there must be some other rational norm that plausibly grounds subjective deontic statuses in all decisions under ignorance, including for all values of n in Table 2.9. Per the methodology I adopted in section 2.4.1, this norm must be one of the other three candidate norms. Neither the MAPV norm nor Medianmax can fill this role, because those norms can ground subjective deontic statuses only when the relevant value is interval-measurable (see above, footnote 57).⁷¹ Accordingly, they cannot ground subjective deontic statuses in all decisions under ignorance.

Nor can Maximax plausibly ground subjective deontic statuses in all decisions under ignorance, because it is possible to construct a counterexample to Maximax that is analogous to Table 2.9:

⁷¹ Moreover, the MAPV norm and Medianmax are not especially plausible when applied to Table 2.9. In that decision, they yield counterintuitive results: (1) if $n < (\frac{1}{2} + \frac{1}{2}\sqrt{5})$, B is obligatory; (2) if $n = (\frac{1}{2} + \frac{1}{2}\sqrt{5})$, either alternative is permissible; and (3) if $n > (\frac{1}{2} + \frac{1}{2}\sqrt{5})$, A is obligatory. It seems arbitrary and counterintuitive for the deontic statuses in Table 2.9 to hinge on this particular irrational number. Of course, seeming arbitrariness and counterintuitive precision can result from the MEV norm as well. But in that case, these can be justified in terms of representation theorems and long-run arguments. Those justificatory bases are not available for the MAPV norm and Medianmax.

Table 2.10

	state 1	state 2
<i>C</i>	$1/n$	$-n$
<i>D</i>	0	0

Surely, we might argue, there is some positive integer, n —make it as high as you like—such that it is (subjectively rationally) impermissible to give up the sure thing guaranteed by Maximin.

That is, there is some point at which (1) D 's minimum possible value of 0 (which is guaranteed by Maximin) is sufficiently higher than C 's minimum possible of $-n$, and (2) C 's maximum possible value of $1/n$ represents a negligible enough gain over Maximin's guaranteed 0, that (3) it is (subjectively rationally) impermissible to take on risk by choosing C . But if Maximax grounds the alternatives' subjective deontic statuses, then it is obligatory to choose C , regardless of the value of n . Because C cannot be both obligatory and impermissible, there are at least some values of n such that Maximax does not ground the alternatives' subjective deontic statuses in Table 2.10. *A fortiori*, Maximax does not ground alternatives' subjective deontic statuses in all decisions under ignorance. Because the MAPV norm and Medianmax cannot do so either, the Expanded n -Value Objection fails.

Next, let's turn to the n -Value Objection. This objection makes the narrower claim that there are *some* decisions under ignorance for which the grounding norm is not Maximin. Equivalently, it makes the claim that there is a subclass of decisions under ignorance in which it is (subjectively rationally) permissible for agents to take on risk. There are, in other words, some decisions under ignorance in which taking on risk is justified (and *ipso facto* permissible). Because Maximin is the "default" norm, the n -Value Objection concedes that decisions under ignorance that fall outside its specified subclass are governed by Maximin (absent some other justificatory basis for taking on risk). This last point, which follows from the assumption that risk is (subjectively rationally) permissible only if justified, is worth emphasizing. Even if the n -

Value Objection is sound, Maximin remains the operative norm for decisions under ignorance that fall outside the objection's specified subclass.

As formulated above, the n -Value Objection specifies a very narrow subclass of decisions under ignorance. Not only must the maximum possible value greatly exceed Maximin's sure thing, but the possible loss from taking on risk must be vanishingly small. The latter attribute was presumably included to make as intuitively strong a case as possible, or perhaps to ensure agreement with the intuitions of as large a portion of the intended audience as possible. But this comes at the cost of decreased relevance. In effect, Table 2.9 is a victim of its own cleverness.

To make as charitable a case as possible, I will consider instead a *Liberalized n-Value Objection*. The liberalized version does away with any specific constraint on possible loss. It holds that a maximum possible value that greatly exceeds Maximin's guaranteed value justifies taking on risk. This version liberalizes the n -Value Objection to make it relevant to as large a subclass of decisions under ignorance as possible.

But even in its liberalized version, the n -Value Objection remains a narrow exception, applicable only in special circumstances. Worse, it applies only to decisions in which the relevant value is interval-measurable. When the relevant value is merely ordinal-measurable, the objection loses its force.

Table 2.11

	state 1	state 2
<i>A</i>	0	100
<i>B</i>	1	2

Table 2.12

	state 1	state 2
<i>A</i>	0	2.1
<i>B</i>	1	2

In Table 2.11, for example, the possible gain under state 2 might be thought large enough to justify taking on risk and choosing *A*. But if the relevant value is merely ordinal-measurable, then Table 2.12—which preserves Table 2.11's ordinal rankings—represents the same state of affairs.

And we need not stop there. Any arbitrarily small difference in value between A's and B's values in state 2—make A's value 2.01, or 2.001, or 2.0001, and so on—will preserve Table 2.11's ordinal rankings, and thus represent the same state of affairs. The maximum possible value can always be reinterpreted as an arbitrarily small gain over Maximin's guaranteed value while preserving the original ordinal rankings. Thus, if the relevant value is merely ordinal-measurable, the notion of a maximum possible value that "greatly exceeds" Maximin's guaranteed value is meaningless. "Greatly exceeds" connotes information about the magnitude of difference between the minimum possible value and the maximum possible value. But ordinal scales convey no information about magnitudes. So, if the relevant value is merely ordinal-measurable, the *n*-Value Objection does not run.

The subclass of EI cases in which the relevant value is ordinal-measurable may turn out to be large, and to include important, non-trivial decisions. Joseph Mendola (1990), for example, argues that the moral value relevant to utilitarianism does not admit of interval-scale measurement. If Mendola is right, or—to generalize—if the relevant value in moral choices where consequentialist considerations are determinative is ordinal-measurable, then the *n*-Value Objection would be inapplicable to these choices.⁷² Maximin would be the operative norm for these choices, even if the *n*-Value Objection is sound.⁷³

⁷² The same can be said of the principle of insufficient reason and the MEV norm. Even I am wrong about the Equiprobability Objection's decisiveness, and the MEV norm (together with the principle of insufficient reason) *does* ground subjective deontic statuses in EI cases with interval-measurable values, it cannot do so in EI choices where the relevant value is ordinal-measurable (see § 2.3.4). So if the moral value relevant to consequentialist considerations is ordinal-measurable, the MEV norm (together with the principle of insufficient reason) cannot ground alternatives' subjective deontic statuses in EI cases where moral consequentialist considerations (rather than, say, deontological considerations) are determinative. Neither can the MAPV norm or Medianmax, because these norms require interval-measurable values. And neither can Maximax, because the *n*-Value Objection, which purports to justify the taking on of risk that Maximax requires of agents, does not apply to choices with ordinal-measurable values. This

More importantly, it is hard to imagine how *any* EI case might plausibly fall under the subclass of decisions under ignorance specified in the Liberalized n -Value Objection. Consider first a two-alternative EI case. Throughout the chapter, I've been numerically representing the alternatives' values as follows:

Table 2.13

	A is better than B	A is worse than B	A is equally as good as B
A	2	1	k
B	1	2	k

For the Liberalized n -Value Objection to apply, the values would instead need to be such that choosing one alternative involves both a possible gain and a possible loss:

Table 2.14

	A is better than B	A is worse than B	A is equally as good as B
A	$2 + l$	$1 - m$	k
B	1	2	k

The value of m may be any positive real number. So long as $m > 0$, the agent takes on risk in choosing A over B . The value of l is more restricted: it must be sufficiently large that A 's value, under the state " A is better than B ," *greatly exceeds* B 's value. Essentially, the value of l must be sufficiently large to trigger the intuition that the possible gain in choosing A is large enough to justify forgoing Maximin's sure thing and taking on risk, despite the lack of subjective probabilities.

I do not see how the values in Table 2.14 can be compatible with the assumption that A and B are apparently incommensurable. Two bearers appear incommensurable to an agent with respect to some value just in case the agent can neither judge one as better than the other, nor

leaves Maximin as the operative norm in moral consequentialist EI cases (if Mendola is right about the ordinal-measurability of the moral value relevant to consequentialism).

⁷³ This is precisely what Mendola himself has argued in more recent work (2014, 2006, 2005)—namely, that a version of Maximin is the operative norm in moral consequentialism, in part because of the ordinal character of classical utilitarianism's hedonic value.

judge them as equally good, with respect to that value. Nor can the agent even hazard a guess as to which is *more likely to be* better than the other.⁷⁴ Yet Table 2.14 depicts an agent who can somehow judge the following:

1. If *A* turns out to be better than *B*, *A* will be *much* better than *B*—whereas if *B* turns out to be better than *A*, *B* will not be as valuable as *A* would be if *A* were better. In other words, the magnitude of the difference between *A* and *B*, if *A* is better than *B*, is greater than the magnitude of the difference between *A* and *B* if *B* is better than *A*.
2. If *A* turns out to be worse than *B*, it will be worse by at least a marginally greater magnitude than *B* would be worse than *A*, if *B* turns out to be worse than *A*.

Then again, for all the agent knows, *A* and *B* might be equal!

Table 2.14 requires an agent who cannot ordinally rank *A* and *B* to nevertheless make sophisticated interval-scale claims about the possible magnitudes of difference between them. Table 2.14 contemplates an agent who can say: “*A* and *B* appear incommensurable to me. I nonetheless have reason to believe not only that *A* and *B* are commensurable, but that they are commensurable in one of several very specific ways.” I cannot imagine how any EI case could fit this description. Table 2.14 seems incompatible with the phenomenon of apparent incommensurability.

The same is true of cases of complete apparent incommensurability with more than two alternatives, for the same reasons. But what about cases of partial apparent incommensurability? Do any of those fall under the subclass of decisions under ignorance specified in the Liberalized *n*-Value Objection?

Here, too, I think that the answer is no. To see why, let’s return to the National Book Award example in which (1) the panelists perceive *A*, the magical realist novel, as incommensurable with either *B* or *C*, but (2) perceive historical fiction novel *B* as slightly better than historical fiction novel *C*. Assume again that this is a case of EI rather than OI, and assume

⁷⁴ See my earlier discussion of the Equiprobability Objection, premise (1).

further (for the purposes of the example) that literary merit is interval-measurable. (If literary merit is ordinal-measurable, the Liberalized n -Value Objection loses its force as applied to this choice.) Under Maximin, B is the only (subjectively rationally) permissible alternative. B maximizes minimum possible value, because under any possible state of the world, there is another alternative— C —that is worse. While there are possible states of the world in which A and C are each ranked lowest, there is no possible state of the world in which B is ranked lowest. B is the only alternative for which there is no possible state of the world in which it is the least valuable alternative.

For the Liberalized n -Value Objection to justify choosing A , the panelists must believe the following. If A is better than B , it must be *much* better than B —its value must greatly exceed B 's. If A is worse than B , the magnitude of the difference between A and B must be greater than it would be if B is worse than A —i.e., the panelists would take on risk in choosing A over B . Yet, for all the panelists know, A might be equal to B . The panelists can say so little about the relative values of A and B that they appear wholly incommensurable. The same tensions arise here that arose in the case of complete apparent incommensurability. An agent⁷⁵ who does not even perceive the alternatives as commensurable at all is nonetheless able to make specific interval-scale claims about the alternatives' relative values in various possible states of the world. Moreover, these claims are not symmetrical: the magnitudes differ, depending on which apparently incommensurable alternative is better—*if* either turns out to be better than the other.

Everything that I have said up to this point is consistent with the possibility that the Liberalized n -Value Objection applies to *some* decisions under ignorance—just not to EI cases. I will finish the chapter by briefly stating why I do not think that the Liberalized n -Value

⁷⁵ In this case it is a collective agent, but that fact is an incidental feature of the particular example I've used to illustrate this case. It does not affect my analysis.

Objection justifies taking on risk in *any* decisions under ignorance. This will insulate my thesis that Maximin grounds subjective rightness in all EI cases from any intuitively plausible counterexamples that a clever reader might devise of EI cases that fall under the Liberalized n -Value Objection.

Here is a decision under ignorance that might be thought to fall within the subset of choices specified by the Liberalized n -Value Objection. In this case, the possible gain under state 2 in choosing A might be thought large enough to justify taking on the risk of loss under state 1:

Table 2.15

	state 1	state 2
A	0	100
B	1	2

Now consider two related decisions. The first has an arbitrarily large number of states in which the maximum possible value obtains. The second has an arbitrarily large number of states in which it does not:

Table 2.16

	state 1	state 2	state 3	...	state n
A	0	100	100	...	100
B	1	2	2	...	2

Table 2.17

	state 1	state 2	state 3	...	state n
A	0	100	0	...	0
B	1	2	1	...	1

Finally, consider a decision that excludes altogether the maximum possible value of 100 for alternative A :

Table 2.18

	state 1
A	0
B	1

Because these are decisions under ignorance, the inclusion of additional states in Tables 2.16 and 2.17 does not affect the likelihood that the maximum possible value of 100 obtains. Nor, strictly speaking, does the exclusion of that value make it any more or less likely to obtain in Table 2.18 than in Table 2.15.

The change in status for the proposition “*A* has a value of 100” between Table 2.18 and Tables 2.15 through 2.17 is not one of probability, but of modality. In Table 2.18, it is *impossible* for *A* to have a value of 100; *necessarily*, it is not the case that *A* has a value of 100. In Tables 2.15 through 2.17, it is *possible* for *A* to have a value of 100; there is at least one possible world in which *A* has a value of 100. This is weaker than the claim that *A* is *more likely* to have a value of 100 in Tables 2.15 through 2.17, because there is nothing about those decision matrices that precludes states 2 through *n* from having probabilities of zero.

In Table 2.18, there is no possible world in which *A* has a value of 100. In Tables 2.15 through 2.17, there is at least one possible world in which *A* has a value of 100. But because of the nature of possible worlds, it is possible in Tables 2.15 through 2.17 to describe an infinite number of possible worlds in which *A* has a value of 100, and an infinite number of possible worlds in which *A* does not have a value of 100.⁷⁶ I do not see how bare modal possibility, without probabilities, can justify taking on risk. Maximin, on the other hand, has a superior modal pedigree: necessity. *Necessarily*—i.e., in every possible world—the agent who chooses according to Maximin’s prescriptions obtains the maximal minimum possible value.

Here is where that leaves us. The default norm for subjective rightness in EI cases and other decisions under ignorance is Maximin. Maximin guarantees a sure thing, while the other candidate norms either permit or require agents to take on risk (by forgoing Maximin’s sure

⁷⁶ There is a possible world exactly like this one in which *A* has a value of 100, a possible world in which *A* has a value of 100 but there is one more grain of sand in the universe, and so on.

thing). Hence, for any decision under ignorance, another candidate norm is operative only if the risk that it permits or requires the agent to take on is justified. The family of *n*-Value Objections constitutes the most promising justification for taking on risk in a decision under ignorance. The Liberalized *n*-Value Objection makes the case that there are at least some decisions under ignorance for which Maximin is not the grounding norm for subjective rightness. But the Liberalized *n*-Value Objection, even if sound, does not apply to any EI cases—nor, in fact, does it justify taking on risk in *any* decisions under ignorance. The Extended *n*-Value Objection, on the other hand, does apply to EI cases, because it applies to all decisions under ignorance. But, as I argued earlier, the Extended *n*-Value Objection is unsound. In the absence of justification for taking on risk we should conclude that Maximin, rather than one of the other three candidate norms, grounds subjective rightness in EI cases.

At long last, we have a full answer to the question of what grounds alternatives' deontic statuses in EI cases (absent deontic constraints). Their objective deontic statuses are grounded in (1) the MAV norm and (2) the fact of the alternatives' actual comparative value. Their subjective deontic statuses are grounded in (1) Maximin and (2) the normative fact of which alternative(s) maximize(s) minimum possible value.

Chapter 3: Apparent Incommensurability

3.1 – OI-EI Uncertainty

3.1.1 – Introduction

In Chapter 1, I examined what grounds alternatives' deontic statuses in OI cases. In Chapter 2, I did the same for EI cases. EI cases involve the added wrinkle of agential ignorance: the agent is ignorant as to the normative fact of which value relation—better than, worse than, or equally good—holds between her alternatives. When a choice scenario involves agential ignorance about relevant normative or empirical facts, it becomes necessary to distinguish between objective and subjective deontic statuses. Although objective deontic statuses may be theoretically important, subjective deontic statuses have far greater practical importance. It is subjective deontic statuses that provide guidance for agents' decisions and ground ascriptions of praiseworthiness or blameworthiness. It is subjective deontic statuses that we invoke when we ask, from the first-personal perspective, "What should I do?" And it is subjective deontic statuses that we invoke when we hold people to account for having done the wrong thing, or praise them for having done the right thing.

In Chapter 3, I further complicate things by introducing another form of agential ignorance. Up to this point, I have assumed that the agent is aware what sort of incommensurability—ontological or epistemic—she faces. But this is not always the case. Agents are sometimes ignorant as to whether a particular case of apparent incommensurability is one of OI or EI. (Call this "OI-EI uncertainty.") Despite its plausibility and (I will argue) ubiquity, OI-EI uncertainty is woefully under-discussed by philosophers. In fact, it does not seem to be discussed at all, at least in print. I am unaware of a single published, sustained treatment of this topic.

OI-EI uncertainty, like the form of agential ignorance at issue in my Chapter 2 discussion of EI cases, is a form of *normative* ignorance or uncertainty. The agent is ignorant or uncertain of the normative fact of whether the apparent incommensurability is ontological or epistemic. But OI-EI uncertainty differs in at least one important respect from the ignorance at issue in Chapter 2. It admits, at least sometimes, of subjective probabilities. Here's an example. I might be fairly confident that two apparently incommensurable political candidates are ontologically commensurable, but believe that the relevant political value is simply too complex—involving too many constituent values interacting in complex ways—for me to discern which candidate is better, or whether they are equally good. (In other words, I am fairly confident that I face an EI case.) Yet I might also remain open to the possibility that they are ontologically *incommensurable* after all. We might (roughly) characterize me as having a subjective probability of 0.8 that my ballot choice is an EI case, and a subjective probability of 0.2 that my ballot choice is an OI case. Other times, my OI-EI uncertainty might be best characterized by an absence of subjective probabilities (i.e., as a decision under ignorance), as was the case for EI-case ignorance in Chapter 2. As we saw in Chapter 2, the term “ignorance” sometimes connotes an absence of subjective probabilities. It is for this reason that I will adopt the term “OI-EI uncertainty,” rather than “OI-EI ignorance.”

In this chapter, I turn at last to the question of how OI-EI uncertainty affects alternatives' deontic statuses. More specifically, I will explore how OI-EI uncertainty affects alternatives' *subjective* deontic statuses. For just as agential ignorance about which value relation obtained in EI cases had no effect on alternatives' *objective* deontic statuses in Chapter 2, so too this new form of agential ignorance does not affect alternatives' objective deontic statuses. Objective deontic statuses are grounded in the agent's actual circumstances, and are not relativized to

agents' beliefs. So, for example, if an agent faces an EI case believing that it *might* be an OI case, her alternatives' objective deontic statuses will nevertheless be grounded in (1) the MAV norm⁷⁷ (since that norm grounds objective deontic statuses in both OI cases and EI cases) and (2) the fact of which alternative maximizes actual value. Her OI-EI uncertainty does not affect the grounding. Or suppose that, on further deliberation, she (mistakenly) becomes certain that she faces an OI case.⁷⁸ In either event—whether she is OI-EI uncertain or mistakenly believes she faces an OI case—it will still be the case that the objectively right alternative(s) is (or are) the one(s) that is (or are) in fact the best, with respect to the relevant value—i.e., the one(s) that maximize(s) actual value. This makes sense, given the role of objective rightness in practical rationality. If she chooses another alternative she will have chosen an objectively wrong alternative, even if she was subjectively rationally permitted to do so.

Of course, the *fact* of whether the case is one of OI or EI may have a dramatic effect on the alternatives' objective deontic statuses—especially for philosophers who theorize that OI cases are governed by an optimizing (rather than maximizing) norm.⁷⁹ But an agent's OI-EI uncertainty—her *beliefs* with respect to that fact—do not affect the alternatives' objective deontic statuses.

⁷⁷ The MAV norm, discussed in Chapter 2, holds that an alternative is rationally permissible (absent deontic constraints) if and only if it maximizes actual (as opposed to, e.g., expected) value.

⁷⁸ Equivalently: she has a subjective probability (or credence) of 1 that she faces an OI case.

⁷⁹ Under an optimizing theory of OI, an alternative is objectively rationally permissible if and only if it is *at least as good as* each other alternative. Accordingly, no member of a set of ontologically incommensurable alternatives is objectively rationally permissible. But if the alternatives were instead *epistemically* incommensurable (i.e., if the agent instead faced an EI case), at least one of the alternatives—namely whichever one(s) maximized actual value—would be objectively rationally permissible. Hence, the same alternative may have different objective deontic statuses, depending on whether the case is one of OI or EI.

How, if at all, does an agent's OI-EI uncertainty affect the *subjective* deontic statuses of apparently incommensurable alternatives? This is a high-stakes question. In section 3.1.2, I argue that in practice, it is difficult—if not impossible—to determine whether any particular apparent incommensurability is case of OI or EI. Before we can claim to have a full account of what grounds subjective deontic statuses in cases of apparent incommensurability, we therefore need a satisfactory account of how and whether OI-EI uncertainty affects subjective deontic statuses.

In section 3.2, I introduce the notion of *stringency* and argue that OI-EI uncertainty affects subjective deontic statuses via the *stringency norm*. Under the soft comparativism framework, subjective deontic statuses are grounded in a norm-belief pair: a rational norm, together with the belief whose propositional content that norm takes as its input. When an agent is OI-EI uncertain, there are two candidate norm-belief pairs: an OI norm-belief pair and an EI norm-belief pair. The agent is uncertain between the two beliefs,⁸⁰ and the properties of OI make the two norms unusually difficult to compare. (One cannot, for example, meaningfully compare both norms using a decision matrix.) The stringency norm is a second-order rational norm which holds that in these sorts of circumstances, a norm-belief pair grounds alternatives' subjective deontic statuses if it is *at least as stringent as* the other candidate norm-belief pair.

In section 3.3, I examine what effect OI-EI uncertainty has on alternatives' subjective deontic statuses under each of the three theories of OI from Chapter 1: maximization, optimization, and parity. Under all three theories, the presence of OI-EI uncertainty in a putative OI case can sometimes affect alternatives' subjective deontic statuses. Section 3.4 concludes.

⁸⁰ That is, she has nonzero subjective probabilities that the propositional content of each belief is true. She has (1) a nonzero subjective probability that she faces an OI case, and (2) a nonzero subjective probability that she faces an EI case.

As I did with the form of ignorance that figured into EI cases in Chapter 2 (see § 2.1.2), I will again restrict my study to cases in which the agent's OI-EI uncertainty is not due to empirical ignorance, time constraints, incompetence, or bad faith.

3.1.2 – It Is Difficult (Perhaps Impossible) To Distinguish OI and EI Cases

In this section, I make the case that most of the time it is difficult to distinguish OI cases from EI cases. It follows that most cases of apparent incommensurability involve OI-EI uncertainty, and it is therefore incumbent on us to determine how OI-EI uncertainty affects alternatives' subjective deontic statuses.

As I said in Chapter 2 (see § 2.2, n.15; § 2.3.2), I wish to remain neutral on the question of whether subjective rightness is a function of the agent's actual beliefs, or of the beliefs that it would be reasonable for her to hold. I will therefore argue both that (1) as a descriptive matter, OI-EI uncertainty is a doxastic attitude that agents have in most cases of apparent incommensurability (i.e., it is an actual belief), and that (2) as a normative matter, OI-EI uncertainty is a doxastic attitude that the agent ought to have (i.e., it is a reasonable belief). Hence, under either account of subjective rightness, OI-EI uncertainty is present in cases of apparent incommensurability in a way that has the potential to affect subjective deontic statuses.

My methodology will be to present a series of high-profile examples (i.e., examples that have been especially influential or are widely cited) that have been offered by others as intuitively clear cases of OI. In each case, I argue that we are not warranted in concluding with certainty that it is an instance of OI. But neither are we warranted in concluding with certainty that it is a case of EI. Instead, the appropriate response is to adopt some degree of OI-EI uncertainty (or alternatively, to adopt a form of OI-EI uncertainty that does not admit of probabilistic representation). Moreover, agents' *actual* responses are typically to adopt at least a

vanishingly small “*de minimis*” subjective probability that they face an EI, rather than OI, case. (And much of the time, the probability will be higher than this.) As we will see later, even a *de minimis* degree of OI-EI uncertainty can be enough to affect alternatives’ subjective deontic statuses.

I begin with the most-cited example⁸¹ of apparent incommensurability in the literature on value incommensurability:

[T]ake as our example the case of a person who has to choose between two options. The one will irrevocably commit him to a career in law, the other will irrevocably commit him to a career as a clarinetist. He is equally suited for both, and he stands an equal chance of success in both. (Raz 1986, 332)

Let’s call the relevant value “goodness of career” (cf. Chang 2002b, 667).

Here, as is most often the case, the apparent incommensurability arises from our attempts to compare two *qualitatively dissimilar* items with respect to a single, overarching value. Both potential careers—one in law, the other in music—instantiate the relevant value (goodness of career) to one degree or another, but they do so in qualitatively disparate ways. The career in law likely pays substantially more, and promises intellectual rigor. The clarinetist career involves the exercise of creativity and artistic expression. Both have opportunities for prestige and glamour, but the prestige of playing in a symphony hall is very different than the prestige of arguing a case before an appellate judge. Both offer the person a means of making the world a better place. But safeguarding the civil rights of the oppressed and engaging in public artistic expression do so in such different ways that it seems impossible to rank which is better with respect to making the world a better place—let alone discern how instantiation of this constituent value (making the world a better place) contributes to instantiation of the overall value (goodness of career).

⁸¹ If an entry in my bibliography has “incommensurability” or “incomparability” in the title, it probably references this example.

Raz tacitly takes all of this to be evidence that the two careers are ontologically incommensurable. “It hardly needs arguing,” writes Raz, “that ... they are incommensurable” (1986, 332). And even if we stipulate that neither is better than the other—i.e., that the only two possibilities are that the careers are equally good or ontologically incommensurable—“[t]he suggestion that they are of exactly the same value cannot be entertained seriously” (1986, 332). Elizabeth Anderson (1993, 57) describes an analogous choice between choosing to live out one’s life in a close-knit, parochial community or leaving to pursue a cosmopolitan life as a ballet dancer, and is similarly dismissive of the possibility of EI in such a case. One “cannot plausibly argue,” writes Anderson, that “the value of ... leading one way of life is commensurable with the value of ... leading the other” (1993, 57).⁸²

Of course, it is plainly true that at least some careers in law are commensurable with some careers in music. A career as a U.S. Supreme Court justice is better than a career as a tone-deaf street busker. Similarly, with respect to Anderson’s example, a life as a respected member of a small rural community is better than life as a homeless city-dweller. But the contention shared by Raz, Anderson, and others is that when the magnitude of difference between two alternatives’ values is less extreme than in those sorts of nominal-notable comparisons,⁸³ qualitative dissimilarity sometimes precludes ontological commensurability.

I agree that these sorts of cases convincingly raise the specter of OI. But I do not see how the difficulty, or even impossibility, of determining how two items relate with respect to some value gives us license to conclude, with certainty, that those particular items are ontologically

⁸² Ruth Chang (2017, 15) similarly suggests that “the significant qualitative difference between [careers] ... make[s] it implausible to think that they are equally good.” Chang takes this to be evidence of parity—more on this below.

⁸³ The term “nominal-notable comparison,” which is now commonly used to describe this sort of comparison, was coined by Ruth Chang (1997b).

incommensurable with respect to that value. My response to Raz's case and examples like it is to remain open to the possibility that the alternatives' apparent incommensurability is merely epistemic. This is different from the position, defended by Donald Regan (1997) and Chris Kelly (2008), that such cases *are* EI cases. My position is instead that the appropriate response to these examples is to adopt some degree of OI-EI uncertainty.

This position is rooted in the notion that normativity is complex. Its component parts—values, reasons, and the like—often behave in erratic, nonlinear ways. Indeed, we need not appeal to cases of apparent incommensurability to see this; it is manifestly apparent in choices whose alternatives are uncontroversially commensurable.

Consider, for example, the ways in which constituent values contribute to an overarching value's structure. Some constituent values—some considerations—are more important than others. For example, when comparing two novels with respect to the value of literary merit, insightfulness plausibly matters more than correct grammar. Or when comparing restaurant menu items, taste plausibly matters more than visual presentation. The fact that some constituent values have greater relative importance is often glossed in terms of the overarching value's weighing the relative strength of each constituent value.

But the analogy of weighing is overly simplistic. For one thing, context matters. Instantiating a particular constituent value to a certain degree may increase an alternative's overall value when combined with one set of instantiations of constituent values, but not when combined with others. (Instantiating poor grammar often decreases a novel's literary merit, but when the poor grammar facilitates accurate depiction of vernacular dialects in a novel whose central literary themes are better served by use of those dialects—as in Mark Twain's *Adventures of Huckleberry Finn*—the poor grammar increases overall literary merit.) Moreover, an

alternative that instantiates to a small degree a constituent value that seems particularly weighty with respect to the overarching value may matter less to the comparison than an instantiation to a large degree of a constituent value that seems less weighty. Visual presentation typically matters less than taste, but spun and molded sugar sculptures of the sort made by master French pastry chefs are so beautiful that their inclusion in a dish might matter more than its insipid flavor.

The analogy of weighing incorrectly suggests that the effects of combining reasons, values, considerations, and the other components of normativity are additive, in the same way that adding a 2-kilogram weight to a 3-kilogram weight yields 5 kilograms. But a better analog is nonlinearity in pharmaceutical drug interactions. Sometimes when we combine chemical compounds, the effects are additive. Other times the effects are synergistic, producing an effect that is greater than the sum of its parts—i.e., greater than the effects would be if additive. And sometimes when particular chemical compounds are combined, the effects may be antagonistic, producing an effect that is less than the effects would be if additive. In fact, the complexity of chemical compound interactions is even more multi-faceted than this partial list suggests. For instance, some compounds, beneficial or neutral on their own, become toxic when combined. And so on.

The components of normativity are similarly complex, even when we set aside deontic constraints. The *Huckleberry Finn* and French pastry examples illustrate two types of non-additive normative phenomena, but there are myriad others. Sometimes, for example, instantiation of a constituent value contributes to overall value only up to a threshold, but instantiation of the constituent value *above* the threshold detracts from overall value. In other words, too much of a good thing becomes a bad thing. Confidence contributes to one's effectiveness as a teacher up to a point, but too much is detrimental. You might become "tone

deaf” and unable to “read the room.” Or you might become so confident that you are no longer able to discern when you don’t know what you’re talking about.⁸⁴ Similarly, a judge’s ability to empathize with parties can make her a better judge—but too much might make her a worse judge by hurting her ability to make impartial decisions. Moreover, the threshold may itself have vague boundaries. Or where the threshold lies may depend on which other constituent values an item instantiates, and to what degree.

Cataloging and modeling non-additive effects is notoriously difficult even in the empirical realm.⁸⁵ It is bound to be harder in the normative realm, where empirical observation is unavailable. Given the complexity of normativity and the impossibility of direct observation, it would be unsurprising if we were sometimes ignorant about certain aspects of a relevant value’s structure. Indeed, normative omniscience seems scarcely more plausible than empirical omniscience. We may not know, or may know only imprecisely, how a particular set of instantiations of constituent values affects the degree to which an alternative, *A*, instantiates the overarching value. And when we must then compare *A* with another alternative, *B*, the result—particularly if *A* and *B* are qualitatively dissimilar—is sometimes that *A* and *B* appear incommensurable.

This is precisely the phenomenon at play in Raz’s career comparison example. One consequence of the alternatives’ qualitative dissimilarity is that they instantiate different constituent values from one another (e.g., good pay and intellectual rigor vs. creativity and artistic expression). This contributes to the appearance of incommensurability because it is not clear in this case how the bearers’ instantiations of particular constituent values contribute to

⁸⁴ Psychologists refer to this as the “Dunning-Kruger effect,” a reference to Dunning & Kruger 1999.

⁸⁵ For a discussion of the challenges of cataloging and modeling the non-additive effects of drug compound combinations, see Bulusu et al. (2016).

their instantiation of the overarching value relevant to the choice. Another consequence of the alternatives' qualitative dissimilarity is that there are some constituent values (e.g., making the world a better place) that the alternatives instantiate in qualitatively disparate ways. Here there is apparent incommensurability with respect to the constituent value, and this contributes to the appearance of incommensurability with respect to the overarching value. And there are still other constituent values (e.g., remuneration) with respect to which the alternatives are commensurable. But commensurability with respect to *some* constituent values does not guarantee commensurability with respect to the overarching value.

Because of the complexity of the overarching value relevant to the choice, and because of the alternatives' qualitative dissimilarities, the alternatives appear incommensurable. Because it is plausible to suppose that we are sometimes ignorant about the structure of the relevant value (particularly in comparisons that involve qualitatively dissimilar alternatives), we should remain open (contra Raz & Anderson) to the possibility that the alternatives are commensurable, and that we are simply ignorant as to which value relation holds between them. But for the same reason, we should remain open (contra Regan & Kelly) to the possibility that the alternatives are genuinely—that is, ontologically—incommensurable.

Donald Regan articulates this position eloquently:

[T]he incomparabilist⁸⁶ is ... wrong in supposing that incomparabilities are ordinarily easy to identify. ... Even if we assume there are some cases of genuine incomparability, it seems very likely that an agent might work long and hard at trying to figure out what her life as a clarinetist and her life as a lawyer would be like and how the values would compare; and she might still find, after she had spent as much time and effort as seemed reasonable even for such an important choice, that she could not be entirely confident *either* that one life was the more valuable *or* that the lives were equally valuable, *or that the lives were incomparable*. (1997, 138-139)

⁸⁶ Regan uses the term “incomparability” to refer to what I call *incommensurability*. See Regan 1997, 129.

Regan goes on in that paper to argue that agents may ignore the last possibility, because his own substantive views preclude the existence of OI cases. He believes that there is one fully commensurable value—G.E. Moore’s notion of “good”—in terms of which all comparisons are made for the purposes of practical rationality. But absent a substantive axiology that disallows ontological incommensurability, the natural and appropriate response is, as the above passage suggests, OI-EI uncertainty.

My claim here regarding OI-EI uncertainty has both a descriptive and a normative component. As a descriptive matter, OI-EI uncertainty is a *natural response* to cases of apparent incommensurability. It is a doxastic attitude that many of us have, much of the time, when confronted with such choices. As a normative matter, some degree of OI-EI uncertainty (or alternatively, OI-EI uncertainty that does not admit of probabilistic representation) is the *appropriate response* to cases of apparent incommensurability. It is a doxastic attitude that I ought to adopt when confronted with such choices.

As I said in Chapter 2, it is a matter of contention whether subjective deontic statuses are grounded in agents’ actual beliefs or in the beliefs that it would be reasonable for an agent to have. If (1) subjective deontic statuses are grounded in the latter, and (2) OI-EI uncertainty is a rationally obligatory doxastic response to apparent incommensurability, then (3) OI-EI uncertainty has the potential to affect alternatives’ subjective deontic statuses in virtually all cases of apparent incommensurability. It is therefore essential to determine whether and how OI-EI uncertainty affects those subjective deontic statuses. If subjective deontic statuses are instead grounded in agents’ actual beliefs, then OI-EI uncertainty still at least has potential relevance to subjective deontic statuses in those cases of apparent incommensurability in which the agent is in fact OI-EI uncertain. So whatever the doxastic ground of subjective deontic statuses—whether

actual or reasonable beliefs—we at least need to determine OI-EI uncertainty's affect on subjective deontic statuses in apparent incommensurability cases in which the agent is actually OI-EI uncertain.

What about apparent incommensurability cases in which the agent does not seem to be OI-EI uncertain? Suppose that Joseph Raz and Donald Regan are contemplating mid-life career changes, each certain that his alternatives are ontologically and epistemically incommensurable, respectively. If you hold the view that *actual beliefs* ground subjective deontic statuses, you might think that we at least needn't worry about OI-EI uncertainty's effect on subjective deontic statuses in cases with OI- or EI-certainty. And if it further turns out to be true that these sorts of cases represent a good proportion of apparent incommensurability cases, then you might accuse me of overstating the breadth of the OI-EI uncertainty problem.

But the case can be made that OI-EI uncertainty *still* has the potential to affect subjective deontic statuses, even if actual beliefs ground subjective deontic statuses and the agent appears to lack OI-EI uncertainty. This is so, firstly, because we must be careful to distinguish cases of certainty (in which the agent has a credence of 1 that, say, her alternatives are OI) from cases of high confidence (in which the agent has a credence that is close to, but less than, 1 that her alternatives are OI). Many cases of apparent incommensurability that look like the former might turn out, on closer inspection, to be the latter. Second, most of us maintain at least a *de minimis* credence—i.e., a small but nonzero subjective probability—that our beliefs are false. My subjective probability that ghosts exist may be vanishingly small, but it is not quite zero. So even in apparent incommensurability cases in which we are highly confident that our alternatives are OI rather than EI (or vice versa), we nonetheless typically retain a *de minimis* credence—make the probability as minuscule as you like—that the alternatives are EI (or vice versa). And, as

we'll see below, even a *de minimis* OI-EI uncertainty can affect alternatives' subjective deontic statuses.

One idea that has been gaining traction in recent years is that some ontologically incommensurable alternatives are nonetheless “comparable” via the nonstandard value relation of parity. Parity's proponents do not overtly deny the possibility that some parity cases might be a kind of EI case, in which the agent is ignorant of the fact that her alternatives are on a par. But they typically assume that parity cases are easily identifiable as such. Ruth Chang (2017) goes so far as to argue that cases of apparent incommensurability, which she calls “hard choices,” just are cases of parity. If this were so, then OI-EI uncertainty would not be the appropriate doxastic response to cases of apparent incommensurability. An agent who faced a case of apparent incommensurability could be confident that she faced a parity case, which is a subtype of OI case.

In her seminal paper “The Possibility of Parity” (2002b), Chang acknowledges that, like those who interpret particular instances of apparent incommensurability as non-parity OI cases, she owes us an explanation as to why alleged cases of parity are not instead EI cases (668-669). Her response is that cases of apparent incommensurability cannot be interpreted as EI cases without adopting an implausible error theory (2002b, 670; see also 1997b, 25). She presents the following example:

Suppose you are a member of a philosophy appointments committee whose task is to compare Aye and Bea for the vacant chair advertised in your department. Aware that predicting potential future merit is difficult, you and your fellow committee members agree that you must determine which candidate has displayed the greatest philosophical talent in her writings to date. Now imagine that you and your fellows have researched both candidates thoroughly; you have examined, discussed, and argued about their written work in great detail, canvassed considered opinions from around the world, dissected letters of recommendation, and so on. Surely it is possible that after careful, cool-headed deliberation, you and the people whose judgment you respect rationally conclude that Aye's work

shows neither more nor less philosophical talent than Bea's. You are not uncertain as to who has the better record to date; rather, the information you have, deemed complete for the purposes at hand, and the care with which you have evaluated it provides positive evidence, not outweighed by any other evidence, that neither is better than the other. (2002b, 670)

Here, a collective agent (a hiring committee) faces a choice between two alternatives (Aye and Bea); the value relevant to the choice is philosophical talent. Each member of the committee independently concludes that none of the standard value relations (better than, worse than, or equally good) holds between the alternatives.⁸⁷ Thus, in order to interpret this as an EI case, we must adopt an error theory about the case, and maintain that the committee members have all judged wrongly.

Chang argues that an error theory about this case is implausible, for several reasons. Because of your thorough research and discussion, you and your colleagues know (virtually) all of the relevant facts. That is, you know everything there is to know about the degree to which Aye and Bea instantiate each of the constituent values. Your judgments are also “systematic,” and “the greater the systematicity of such judgments, the less reason there is to think that the error theory is correct” (2002b, 671). Moreover, you all agree—you each form the same judgment that Aye is neither more, less, nor equal in philosophical talent with Bea. And finally, according to Chang, it is plausible that this sort of case is common. The error theorist “owes us an explanation of where we go wrong,” and “of why we systematically make this putative error” (2002b, 671).

What sort of error theory can a proponent of the EI interpretation offer, and is it as implausible as Chang suggests? Your error in this case cannot arise from ignorance of any

⁸⁷ Although Chang's description of the example is technically consistent with committee members judging that Aye and Bea are equally good (the members judge that “neither is better than the other”), it is contextually obvious that Chang intends her example to exclude this possibility.

relevant empirical facts. Nor does it seem particularly plausible that your error arises from any ignorance about the relevant, overarching value's content. As professional philosophers, you all have a good grasp of which constituent values are relevant to philosophical talent. You all know that originality, ability to identify and respond to powerful objections, etc. are relevant, while sexual orientation and cooking skills are not. The most plausible error theory that the EI proponent can offer is that you are partially ignorant with respect to the relevant value's *structure*. That is, there is a feature of the overarching value of philosophical talent that determines how Aye's particular instantiations of the constituent values affect her overall value relative to Bea. You and your colleagues are ignorant of what the overarching value's structure determines in this case. Although you know all relevant empirical facts, you are ignorant of this normative fact.

It is not surprising, the EI proponent can argue, that you and your colleagues are ignorant of what the overarching value's structure determines in this case. After all, the way in which an alternative's set of instantiations of constituent values determines its overall value is, as we just discussed, a complex business. It eludes straightforward modeling via weighting of constituent values. And each manifestation of a relevant attribute is context sensitive; how it affects an alternative's overall value is dependent on how that alternative manifests other constituent values. Given the complexity of this phenomenon, the EI proponent can contend, it is perfectly plausible that you are ignorant as to some of its finer points.

In Chang's scenario, however, the EI proponent must account for a further feature: you and your colleagues all agree. You all conclude that Aye is neither better than, nor worse than, nor equal to Bea. Chang might reply that this group consensus makes the error theory less plausible. Note, though, that this reply is compelling only if, as Chang further alleges, consensus

in cases of apparent incommensurability among people who have knowledge of the relevant empirical facts and are experts about the overarching value is common. Isolated cases can be fairly easily dismissed by the error theorist as chance occurrences.

It is certainly plausible that cases of apparent incommensurability are common. It is a controversial empirical question whether *this particular type* of apparent incommensurability case is common. There are a number of factors that are essential to Chang's example. For her error theory argument to be persuasive, these factors must all be present together in a large number of cases. First, there must be a group of people evaluating alternatives with respect to some value, all of whom are *experts* on that value. Otherwise, they may be ignorant of important features of the relevant value. Chang's case is compelling not only because the evaluators agree, but also because they know what they're doing. This is important to minimize the possibility of psychological phenomena like groupthink, which those who favor the EI-case interpretation might otherwise propose as an alternative causal explanation for consensus.

Second, the experts must know *all* of the relevant empirical facts. No doubt hiring committees typically discharge their duties with appropriate diligence. Nonetheless, it seems likely that in many cases one member may have read the candidates' dossiers less closely than another, and thus have an incomplete knowledge of how the candidates instantiate the constituent values.

Moreover, the experts must make the same evaluative judgment *for the same reasons*. It is not enough that they share the same evidence and reach the same conclusion; they must take their shared body of evidence as supporting the conclusion for the same reasons. If they take their evidence as requiring the same conclusion for different reasons, then they are disagreeing about how the alternatives' instantiations of the constituent values affect their overall value.

They are disagreeing about the value's structure. But disagreement about the value's structure is consistent with an EI-case interpretation of apparent incommensurability, and is thus consistent with the error theory.

Even in Chang's hiring committee case, the EI proponent can provide an explanation of where the ignorance lies. Although the committee members know all relevant empirical facts, they are ignorant of some normative facts. Specifically, they are partially ignorant of what the overarching value's structure determines in this case. In addition, the EI proponent can respond to the consensus in Chang's example by contending that this sort of case is not as widespread as she suggests.

My goal in setting all of this out is not to convince you that Chang's example is an instance of EI rather than parity-type OI. Instead, my more modest aim is to argue that even on the assumption that parity cases exist, they are not nearly so easy to identify as Chang supposes. Putative parity cases, like other putative OI cases, are more controversial than their proponents suggest. Here, too, the natural and appropriate response is some degree of OI-EI uncertainty.

Chang (2002b; 2017) also describes another sort of (alleged) parity case which, at first glance, seems less amenable to OI-EI uncertainty. This subclass of parity cases involves choices in which the agent has "first-personal authority" over the content and structure of the relevant value. Here is an example:

Suppose you must determine which of a cup of coffee and a cup of tea tastes better to you. The coffee has a full-bodied, sharp, pungent taste, and the tea has a warm, soothing, fragrant taste. It is surely possible that you rationally judge that the cup of Sumatra Gold tastes neither better nor worse than the cup of Pearl Jasmine and that although a slightly more fragrant cup of the Jasmine would taste better than the original, the more fragrant Jasmine would not taste better than the cup of coffee. In this case, it is plausible to suppose that you know everything that is relevant to comparing the drinks and that in this case you have first-person authority over which tastes better to you; your judgment about their comparative tastes provides the truth as to their comparative taste to you. On what ground

could the skeptic make out that these judgments are mistaken? If you know everything that is relevant to the judgments, the skeptic would have to insist that whether the one drink tastes better to you or not is in principle unknowable. But surely we can know which of the two drinks tastes better to us in such cases. (2002b, 669)

To interpret this case of apparent incommensurability as an EI case, we would need to judge the agent as ignorant as to whether one beverage tastes better to her, or whether they taste equally good. But the problem with supposing that she is ignorant as to their relative tastiness, says Chang, is that “judgment about their comparative tastes *provides the truth as to their comparative taste* to her” (emphasis added). The more fragrant cup of Jasmine tastes better than the less fragrant Jasmine in virtue of the fact that the agent judges it to taste better. It is better (or worse, etc.) *if and only if, and because*, the agent judges that it is better. Similarly, the coffee tastes neither better nor worse than nor equally as good as the tea if and only if, and because, the agent judges this to be the case. Here, if not our earlier two examples, we surely have an iron-clad case of OI—or so the argument goes.

In response, it is first worth noting that even if Chang’s argument here succeeds, it rules out the possibility of EI only when the agent has “first-personal authority” over the relevant value. This doesn’t necessarily narrow the example’s applicability to a small number of cases—we make judgments about comparative taste all the time—but it might narrow the example’s applicability only to trivial cases. The relevant value in most weighty, important choices includes constituent values over which the agent does not have first-personal authority. Agents do not have first-personal authority, for example, over most moral values.⁸⁸ Chang’s example is also

⁸⁸ The possible exception is the moral (dis)value of the agent’s own pleasure (or pain). But even if the correct moral theory is a staunch utilitarianism according to which the only moral value is pleasure/pain, agents would still not have first-personal authority when it comes to determining the value of *others’* pleasure/pain.

inapplicable to whole domains of practical rationality. For example, judges and lawyers do not have first-personal authority over legal values.

Chang (2017, 5-6; 2002b, 670-671) argues that we can extrapolate from the first-personal authority examples to other cases of apparent incommensurability in which agents lack first-personal authority over the relevant value. But for the reasons already laid out in my discussion of her hiring committee example (which demonstrates how the extrapolation is supposed to work), I do not think that we can do so with certainty. At most, her hiring example discussion succeeds in increasing our subjective probabilities in the proposition that for some particular case, ϕ , that resembles the hiring committee example, ϕ is a case of parity. But so long as we retain some nonzero subjective probability in the proposition that ϕ is an EI case, we retain OI-EI uncertainty.

Second, Chang's account of taste comparisons obscures the fact that judgments of comparative taste can sometimes be hard to make. The relevant value in such judgments is not literal taste—the combination of salty, sweet, sour, bitter, and umami perceived by the taste buds. It is metaphorical taste, of which literal taste is one constituent value. Aroma—the bundle of sense perceptions contributed by the olfactory system—is another constituent value. But those constitute only a partial list. There is also a whole web of Proustian and contextual associations that affect taste. South American coffees might tend to taste better to me than coffees from other parts of the world because their aromas evoke my childhood memories of my grandmother's living room. Or coffee might taste better to me at certain times of day (say, first thing in the morning) than others.

Even a seemingly straightforward value like taste turns out, on closer inspection, to be surprisingly complex. And this complexity partially accounts for why judgments of comparative

taste are sometimes hard to make. The complexity of the relevant, overarching value partially explains the “agony” we sometimes feel when trying to decide which of two things tastes better to us: the hemming and hawing, the holding up of one alternative’s qualitative attributes against the others, and so on. Chang might reply that her account *can* take account of this. The agent agonizes, then judges. And once the agent judges, the propositional content of her judgment is true—by virtue of her having so judged, since taste is a value over which she has first-personal authority. But this account cannot make sense of another reason why we sometimes agonize: we want to get the “right” answer. And wanting to get the right answer is comprehensible only if it is possible for us to get the wrong answer—to make a judgment whose propositional content is false.

Chang’s notion of first-personal authority cannot account for the fact that we sometimes get comparative taste judgments wrong. Suppose we sample pints of a couple of qualitatively dissimilar beers at a local brewery—say, an English porter and a Belgian tripel. Having taken one sip of each, I tell you that while I’ve only taken one sip and haven’t yet given much thought to the matter, I judge that the tripel tastes better. But over the course of our evening I drink both pints, carefully noting their various qualitative attributes as I do. I close my eyes and attend to the particular malty sweetness of the porter and the crisp, dry finish of the tripel. I take in the Porter’s coffee notes and the tripel’s alcoholic warmth, and let their respective Proustian associations wash over me. When it’s time to go, I declare that I was initially mistaken: the stout is actually the better-tasting beer.

If my judgments “provided the truth” (as Chang puts it) as to their propositional content, then I could not have been mistaken. The propositional content of my judgment would be true if and only if, and because, it was the propositional content of my judgment. Necessarily, my initial

judgment that the tripel was better would be correct (as would my subsequent judgment that the stout is better).

Chang's account of the value of taste (and other values over which agents have first-personal authority) entails that the beers' relative values *changed*: originally the tripel was better, but now the stout is better. No doubt this is sometimes what happens. Our tastes can and do change dramatically—often over the course of years, but sometimes within a much shorter timeframe. Nevertheless, it strikes me as plausible that at least sometimes we do conclude that our initial judgment was mistaken. If this is so, then we are forced to conclude, contra Chang, that comparative judgments made with respect to a first-personal-authority value do not ground the truth of their propositional content. Instead, the truth of their content is grounded in the surprisingly complex value of personal taste. A careful evaluation is (as with any evaluative comparison) more likely to yield a true judgment of their relative value than an off-the-cuff one. But even a careful evaluation might leave me flummoxed, facing an apparent incommensurability. At the end of our brewery outing, I might instead conclude that I am no longer confident which beer is better, nor even that they are equally good.

The takeaway here is that even with values such as personal taste, it is still possible to be ignorant of items' relative value. Sometimes further deliberation resolves the apparent incommensurability, but other times it persists. Once again, my aim is not to argue that apparent incommensurabilities in taste comparisons *are* EI cases. I aim only to rebut Chang's assertion that their status as OI cases is unassailable. Even here, some degree of OI-EI uncertainty—make it as small as you like—is the natural and appropriate response.

Finally, I'll reemphasize my earlier point about *de minimis* credences. Even if we believe that Chang is right about first-personal-authority values, we will still likely retain a *de minimis*

credence that we are mistaken about this fact, and that EI cases in which the relevant value is a first-personal-authority value are possible. If this is so, then we will likely also retain a *de minimis* level of OI-EI uncertainty in cases of apparent incommensurability in which the relevant value is one over which we have first-personal authority. And, as we will see, even a *de minimis* level of OI-EI uncertainty can sometimes be enough to affect alternatives' subjective deontic statuses.

In this section, I've been making the case that OI-EI uncertainty is present in virtually all cases of apparent incommensurability. In practice, it is often difficult or impossible to distinguish OI from EI cases, and even when we think we have managed to identify (say) a genuine OI case, some OI-EI uncertainty remains. This second point is true both descriptively and normatively. Descriptively, we are likely to retain at least a *de minimis* degree of OI-EI uncertainty even when we are highly confident that we've identified an OI (or EI) case. And normatively, some level of OI-EI uncertainty (even if only *de minimis*) is typically obligatory.

Because OI-EI uncertainty is a component of agents' actual beliefs and is a rationally required doxastic attitude, it has the potential to affect alternatives' subjective deontic statuses regardless of whether subjective rightness is a function of the agent's actual beliefs or the beliefs that it would be reasonable for her to have. And because OI-EI uncertainty is present in virtually all cases of apparent incommensurability, it is crucial that we examine how, if at all, OI-EI uncertainty affects alternatives' subjective deontic statuses.

3.2 – OI-EI Uncertainty Affects Subjective Deontic Statuses Via the Stringency Norm

In the last section I argued that some degree of OI-EI uncertainty is both descriptively present and normatively required in most cases of apparent incommensurability. In this section I ask how, if at all, this OI-EI uncertainty affects alternatives' subjective deontic statuses. I

propose that we understand this question, like the principal questions in Chapters 1 and 2, as one about the metaphysical grounding of alternatives' deontic statuses. Here is the question, restated in those terms: *in cases of apparent incommensurability that involve OI-EI uncertainty, what grounds alternatives' subjective deontic statuses?*

The simplest approach to modeling OI-EI uncertainty would be to simply add an OI column to each of the decision matrices in Chapter 2. If we add an additional “incommensurable” state to the existing partition, the resulting set of states {better than, worse than, equally good, incommensurable} can still plausibly be taken as mutually exclusive and exhaustive. Moreover, we can reinterpret the EI cases in Chapter 2 as having already included an implicit “incommensurable” state with a subjective probability of zero. To model OI-EI uncertainty, we would simply make this implicit state explicit and assume that it has a nonzero probability. In one respect we would even be better off than in Chapter 2, because depending on the particular choice in question, we might now have at least partial information about the states' subjective probabilities. We might know, for example, that the “incommensurable” state has a probability of 0.7, while the other states jointly have a probability of 0.3. We could then run through the same analyses as we did in Chapter 2, noting whether the new set of numerical values and subjective probabilities affects any of the candidate norms' prescriptions and plausibility.

Unfortunately this approach will not work, because it is impossible to numerically represent ontological incommensurability. Representing alternative *A* with a higher (or lower) number than alternative *B* implies that *A* is better than (or worse than) *B*. And representing *A* and *B* with equal numbers implies that *A* is equally as good as *B*. (In fact, representing OI with numeric equality would be problematic even if we did not need to simultaneously represent equal

goodness elsewhere in the same decision matrix. Equal goodness is plausibly represented by numeric equality because the value relation and the numeric relation share important traits: both are plausibly symmetric, transitive,⁸⁹ and reflexive. But OI does not share all of these features. For example, OI is not reflexive: *A* is not incommensurable with *A*.) Of course, we could add an OI state to the decision matrices in Chapter 2 and represent the values in that column non-numerically—say, with letters instead of numbers. But doing so would not be informative. The resulting matrix would not facilitate meaningful comparisons of alternatives' values across states. It would not provide adequate grounds for assessing the candidate rational norms' plausibility. We will need to find a different approach.

Under the soft comparativism framework I've adopted, deontic statuses are typically grounded (absent deontic constraints) in a norm-fact or norm-belief pair. The former grounds objective deontic statuses; the latter grounds subjective deontic statuses. We have two candidate norm-belief pairs that might ground subjective deontic statuses in a case of apparent incommensurability. One norm-belief pair grounds subjective deontic statuses if the choice is an EI case (viz., Maximin⁹⁰ and the belief that the alternatives are EI). A different norm-belief pair grounds subjective deontic statuses if the choice is an OI case (viz., either maximizing, optimizing, or optimizing + parity⁹¹ depending on your OI-norm theory, together with the belief

⁸⁹ For challenges to the standard assumption that equal goodness is transitive, see Temkin 2012.

⁹⁰ I argued in Chapter 2 that Maximin is the operative rational norm in EI cases.

⁹¹ From Chapter 1: (1) under the maximizing norm, an alternative is rationally permissible if and only if it is *not worse than* each other alternative; (2) under the optimizing norm, an alternative is rationally permissible if and only if it is *at least as good as* each other alternative; and (3) under the optimizing + parity norm, an alternative is rationally permissible if and only if it is *either* at least as good as *or on a par with* each other alternative. As I said in Chapter 1, I will remain neutral on the question of which of these three norms grounds alternatives' deontic statuses in OI cases. In this chapter (Chapter 3), I will examine how OI-EI uncertainty affects alternatives' subjective deontic statuses under each of the three OI-norm theories. That way the usefulness of my work in Chapter 3 is not tied to any particular OI-norm theory.

that the alternatives are OI). Which norm is operative depends on which belief the agent either has or should have. Because the agent is and should be uncertain between the two beliefs (i.e., she is and should be OI-EI uncertain), it is not clear which of the two norms is operative. It is not clear which norm-belief pair grounds alternatives' subjective deontic statuses.

One alternative approach, then, is to see if there is a plausible second-order norm that somehow adjudicates between the two candidate norm-belief pairs. If so, the second-order norm and whichever first-order norm it picks out (Maximin or maximizing/optimizing/optimizing+parity) could jointly ground the alternatives' subjective deontic statuses, together with whichever belief the first-order norm takes as inputs.

This is the approach I will take, and in a moment I will propose the second-order norm that I believe is operative in cases of apparent incommensurability that involve OI-EI uncertainty (which, as I argued in the last section, is most cases of apparent incommensurability). But before I can do that, I need to lay the conceptual groundwork for my proposed norm by introducing the notion of *stringency* of norm-belief pairs. Roughly stated, stringency has to do with the idea that with respect to a particular choice, some norm-belief pairs are stricter than others. That is, they render a larger number of alternatives (subjectively rationally) impermissible. Hence, they do not permit the agent as wide a latitude when choosing between alternatives. That's the general idea, though the technical notion of stringency—which I'll define in the next paragraph—captures a narrower range of phenomena than this rough gloss suggests.

Stringency is a relation that holds between norm-belief pairs relative to a set of alternatives. Given a set of alternatives, Δ , and two norm-belief pairs x and y :

1. **x is more stringent than y** if and only if: (i) there is at least one alternative in Δ that x , but not y , renders (subjectively rationally) impermissible; but (ii) there no alternative in Δ that y , but not x , renders (subjectively rationally) impermissible.

2. ***x* and *y* are equally stringent** if and only if neither renders an alternative impermissible that the other renders permissible.⁹²

Suppose, for example, that an agent must choose between two alternatives: $\{A, B\}$. According to norm-belief pair x , A is (subjectively rationally) obligatory and B is (subjectively rationally) impermissible. But according to norm-belief pair y , A and B are both (subjectively rationally) permissible. On those facts, x is more stringent than y . There is an alternative that x , but not y , renders (subjectively rationally) impermissible: B . But there is no alternative that y , but not x , renders (subjectively rationally) impermissible, because y renders both alternatives (subjectively rationally) permissible.

By obeying x 's prescriptions, the agent simultaneously obeys y 's prescriptions. It is not possible to do otherwise, because when (as here) x is more stringent than y , there is no alternative that is (subjectively rationally) permissible under x but (subjectively rationally) impermissible under y . On the other hand, it is possible to obey y 's prescriptions and simultaneously violate x 's prescriptions—in this case, by choosing B .

Stringency is analogous to a more familiar notion from decision theory: dominance. Dominance is a relation that holds between alternatives (or acts) relative to a set of values (or outcomes). One alternative dominates another if and only if: (1) in each state, its value is equal to or greater than the other alternative's value; and (2) there is at least one state in which its value is greater than the other alternative's.

Table 3.1

	state 1	state 2
<i>C</i>	1	2 ← higher value than <i>D</i>
<i>D</i>	1	1

⁹² It is also possible to discuss stringency in the context of objective (vs. subjective) rightness by (1) substituting norm-fact pairs for the norm-belief pairs, and (2) substituting objective deontic statuses for the subjective deontic statuses.

In Table 3.1, alternative *C* dominates alternative *D*. For visuospatial thinkers, here is an analogous table depicting stringency in the preceding example of the choice between *A* and *B*:

Table 3.2

	<i>A</i>	<i>B</i>
<i>x</i>	permissible	impermissible ← stricter deontic status than under <i>y</i>
<i>y</i>	permissible	permissible

Note that unlike Table 3.1, Table 3.2 is not a standard decision matrix. The rows and columns in a decision matrix depict alternatives and states, respectively; the rows and columns in Table 3.2 instead depict norm-belief pairs and alternatives, respectively. The reason I've placed alternatives in columns instead of rows in Table 3.2 is to create a sense of visual parity with Table 3.1.

While dominance holds between alternatives relative to a set of numeric values, stringency holds between norm-belief pairs relative to a set of deontic statuses. To spell out the analogy: “*x* is more stringent than *y*” is to deontic statuses as “*C* dominates *D*” is to numeric values. In the states in which *C* and *D* have equal values (state 1 in Table 3.1), an agent who chooses *C* will obtain the same amount of value as she would if she chose *D*. And in states in which *C*'s value is greater than *D*'s (state 2), she will obtain more value than she would if she chose *D*. In other words, an agent who chooses *C* cannot obtain less value, and might obtain more value, than if she chose *D*.

Similarly, an agent who acts in accordance with *x*'s deontic statuses cannot satisfy fewer candidate norm-belief pairs' deontic statuses, and might satisfy more candidate norm-belief pairs' deontic statuses, than if she acts in accordance with *y*'s deontic statuses. For those alternatives that are (subjectively rationally) permissible according to both *x* and *y* (e.g., *A* in the Table 3.2), an agent who acts in accordance with *x*'s deontic statuses (e.g., by choosing *A*) will simultaneously satisfy *y*'s deontic statuses. And for those alternatives that are permissible under

y but not x (e.g., B), an agent who acts in accordance with x 's deontic statuses (by instead choosing some other permissible alternative—in this case, A) will simultaneously satisfy y 's deontic statuses (because that other alternative— A —is also permissible under y)—whereas, if she had instead chosen B , she would satisfy y 's deontic statuses but not x 's deontic statuses. In other words, by obeying the stricter of the two norm-belief pairs, the agent simultaneously obeys the laxer norm-belief pair.

The stringency-dominance analogy holds only for the relation “more stringent than.” “Equally stringent” is instead analogous to a set of alternatives whose values in each state are equal (call these “identical-value alternatives”), as in Table 3.3:

Table 3.3: Example of identical-value alternatives

	state 1	state 2
E	1	2
F	1	2

Table 3.4: Example of equally stringent norm-belief pairs

	G	H
x	permissible	impermissible
y	permissible	impermissible

Note that just as Table 3.5 (below) is *not* an example of identical-value alternatives, Table 3.6 is *not* an example of equally stringent norm-belief pairs:

Table 3.5

	state 1	state 2
G	2	1
H	1	2

Table 3.6

	G	H
x	impermissible	permissible
y	permissible	impermissible

I have included Table 3.6 to illustrate the point that “equally stringent” does not mean “yields the same number of impermissible alternatives.” Instead, it means “renders an identical set of alternatives (subjectively rationally) impermissible.” If x and y are equally stringent, then an alternative is (subjectively rationally) impermissible under x if and only if it is (subjectively rationally) impermissible under y .

There are plausible first-order rational norms that take as their inputs (1) the fact or belief that one alternative dominates another and (2) the fact or belief that an agent faces identical-value alternatives. Here is a common formulation of the norm related to dominance:

dominance norm (version 1): Dominated alternatives are rationally impermissible (absent deontic constraints).⁹³

If an agent faces exactly two alternatives, then the dominant alternative becomes, by process of elimination, obligatory. Hence:

dominance norm (version 2): Given a choice between two alternatives $\{A, B\}$, if A dominates B , then B is impermissible and A is obligatory (absent deontic constraints).

Here is the plausible norm that pertains to identical-value alternatives:

IVA norm (version 1): Identical-value alternatives have identical deontic statuses (absent deontic constraints).⁹⁴

It is, of course, possible that deontic constraints might sometimes lead identical-value alternatives to have different deontic statuses. An exclusionary reason, for example, might make it the case that one of two identical-value alternatives is impermissible and the other is obligatory. But absent deontic constraints, it is hard to imagine how two identical-value alternatives could have different deontic statuses.

The IVA norm, like the dominance norm, can be reformulated for choices between exactly two alternatives:

IVA norm (version 2): Given a choice between two alternatives $\{A, B\}$, if A and B are identical-value, then A and B are both rationally permissible (absent deontic constraints).

⁹³ The dominance norm may render either objective or subjective deontic statuses, depending on whether the propositional content that it takes as its input is that of a fact or of a belief.

⁹⁴ Like the dominance norm, the IVA norm may render either objective or subjective deontic statuses, depending on the source of its input.

If an agent faces exactly two alternatives and those alternatives are identical-value, there is nothing to ground a prescription that the agent choose one rather than the other (assuming an absence of deontic constraints). Both are rationally permissible.

I propose that just as dominance and identical-value alternatives figure into plausible first-order rational norms, so too stringency figures into plausible second-order rational norms. Here is the second-order norm that analogizes to the dominance norm (version 2):

“more stringent than” norm: Given two candidate norm-belief pairs $\{x,y\}$ such that the agent is (or should be) uncertain between the two beliefs and the two norms are not directly comparable, if x is more stringent than y , then x grounds the alternatives’ subjective deontic statuses and y does not.

“Uncertainty” here means that the agent has a nonzero subjective probability that the propositional content of each belief is true. It’s worth emphasizing here that it is the *beliefs*, not the *norms*, that the agent is uncertain between. The norms are those that would be operative (i.e., ground alternatives’ subjective deontic statuses), conditioned on the agent’s holding each respective belief. I say “is (or should be) uncertain” to hedge on the question of whether it is actual or reasonable beliefs that ground subjective rightness. (Since OI-EI uncertainty is both descriptively present *and* normatively required in cases of apparent incommensurability, I need not take a stance on this question.) By “not directly comparable,” I mean that the norms cannot be directly compared by means of something like a decision matrix. They are not amenable to the kind of comparative analysis that I used in Chapter 2 to assess candidate rational norms’ plausibility. Thus, the “more stringent than” norm applies only to unusual cases in which: (1) the agent is uncertain between two beliefs; (2) a different rational norm would be operative if the agent held one belief than if she held the other; and (3) those two norms are not directly comparable.

Unlike the dominance norm, which can ground both objective and subjective deontic statuses, the “more stringent than” norm grounds only subjective deontic statuses. This is because the “more stringent than” norm takes agential uncertainty—specifically, uncertainty between the beliefs in each of two norm-belief pairs—as an input. This is an essential feature of the norm because it is the uncertainty between beliefs that (partially) determines which norm-belief pairs are relevant to the choice.

The “more stringent than” norm is plausible for the same sorts of reasons that the dominance norm is plausible. In Table 3.1, the agent could do no worse, and might do better, by choosing *C* instead of *D*. By choosing *C* she can simultaneously acquire all of the value she’d obtain if she chose *D*, and might moreover acquire more value. It is therefore plausible to suppose that she is rationally required to do so. Similarly in Table 3.2, the agent, by obeying *x*’s prescriptions, can simultaneously obey all of *y*’s prescriptions, while moreover obeying *x*’s stricter prescriptions with respect to those alternatives for which *y* yields a laxer subjective deontic status.

This consideration would be irrelevant in a discussion of *objective* rightness. An agent’s uncertainty with respect to which norm-fact pair grounds her alternatives’ objective deontic statuses has no effect on which of those norm-fact pairs is in fact operative, because objective rightness does not take the agent’s uncertainties and ignorance into account. But the “more stringent than” norm applies only to *subjective* rightness. And in that context, this consideration is compelling.

Just as the dominance norm does not take states’ probabilities into account, the “more stringent than” norm does not take the subjective probabilities of the norm-belief pairs’ belief components into account. Suppose (for simplicity) that there is exactly one state, *S*, in which the

dominant alternative is better. It does not matter how unlikely it is that S obtains, because under all other states, the dominant and dominated alternatives have equal values. The inclusion of the dominant state in the partition “tips the scales” in favor of the dominant alternative. The dominant alternative offers the (perhaps minuscule) possibility of obtaining more value than she would obtain if she’d chosen the dominated alternative, without requiring the agent to take on risk. Because she needn’t take on any risk in choosing the dominant alternative, it does not matter whether or not S ’s probability is high or low, so long as it is nonzero.

With regard to the “more stringent than” norm, it similarly does not matter how high or low the subjective probability of the more stringent norm-belief pair’s belief is, so long as it is nonzero. Suppose (for simplicity) that there is exactly one alternative, A that the more stringent norm-belief pair renders (subjectively rationally) impermissible. By obeying the more stringent candidate norm-belief pair’s prescriptions, the agent guarantees that she won’t violate its prescriptions (by choosing A). Moreover, she is able to do so without taking on any risk of violating the less stringent candidate norm-belief pair’s prescriptions: by obeying the more stringent norm-belief pair’s prescriptions, she simultaneously ensures that she obeys the less stringent norm-belief pair’s prescriptions.

Here is the second-order norm that analogizes to the IVA norm (version 2):

“equally stringent” norm: Given two candidate norm-belief pairs $\{x,y\}$ such that the agent is (or should be) uncertain between the two beliefs and the two norms are not directly comparable, if x and y are equally stringent, then both x and y ground the alternatives’ subjective deontic statuses.

Like the “more stringent than” norm, the “equally stringent” norm applies only to unusual cases in which: (1) the agent is uncertain between two beliefs; (2) a different rational norm would be operative if the agent held one belief rather than the other; and (3) those two norms are not directly comparable.

To see why this norm is plausible, consider again Table 3.4:

Tale 3.4 (reprinted)

	<i>G</i>	<i>H</i>
<i>x</i>	permissible	impermissible
<i>y</i>	permissible	impermissible

In this choice, (1) the agent is uncertain which of two candidate norms grounds her alternatives' subjective deontic statuses, and (2) both of those norms, if operative, would render identical subjective deontic statuses anyway. Subjective rightness consists in something being right relative to the agent's beliefs, and the "equally stringent" norm strikes me as a plausible account of how these particular beliefs affect the subjective rightness of her choice. *Because* she is uncertain which norm operates, and *because* both norms would render identical subjective deontic statuses anyway, her alternatives' subjective deontic statuses match those rendered by both norms. By grounding the subjective deontic statuses in *both* norms, the second-order "equally stringent" norm interprets these *because*'s as expressing a metaphysical grounding relation.⁹⁵

Finally, if we take "*x* is at least as stringent as *y*" to mean that *x* is either equally as stringent as or more stringent than *y*, then we can tidily combine both the "more stringent than" norm and the "equally stringent" norm into a single, second-order norm:

stringency norm: Given two candidate norm-belief pairs $\{x,y\}$ such that the agent is (or should be) uncertain between the two beliefs and the two norms are not directly comparable, if *x* is at least as stringent as *y*, then *x* grounds the alternatives' subjective deontic statuses.

The stringency norm, like the dominance norm and IVA norm, has limited application. The dominance and IVA norms offer no advice to an agent facing a choice like the following:

⁹⁵ Recall from Chapter 1 (§ 1.1) that I am remaining neutral between competing accounts of metaphysical grounding.

Table 3.7

	state 1	state 2	state 3
<i>A</i>	5	2	1
<i>B</i>	3	4	1

When (as here) neither alternative dominates the other, and the alternatives are not identical-value, the dominance and IVA norms are inapplicable. And just as not all sets of numeric values are amenable to a dominance or IVA analysis, not all sets of deontic statuses are amenable to a stringency analysis. When neither norm-belief pair is more stringent than the other, and the norm-belief pairs are not equally stringent, the stringency norm is inapplicable. For example:

Table 3.8

	<i>C</i>	<i>D</i>	<i>E</i>
<i>x</i>	impermissible	permissible	permissible
<i>y</i>	permissible	impermissible	permissible

But the important question when considering the usefulness of any of these norms in grounding agents' subjective deontic statuses and in guiding their decisions is not whether they are operative in every choice, but merely whether they are operative in the particular choice being considered.

We now have the tools we need to resolve the problem of OI-EI uncertainty in cases of apparent incommensurability. As a refresher, here is the situation in which we found ourselves at the beginning of this section. We have two candidate norm-belief pairs that might ground subjective deontic statuses in cases of apparent incommensurability. One is the norm-belief pair that would ground subjective deontic statuses if the choice were an OI case. This consists of (1) a maximizing, optimizing, or optimizing + parity norm, depending on your theory of OI, and (2) the belief that the alternatives are ontologically incommensurable. The other is the norm-belief pair that would ground subjective deontic statuses if the choice were an EI case. This consists of (1) Maximin and (2) the belief that the alternatives are epistemically incommensurable. Because

of the OI-EI uncertainty present in cases of apparent incommensurability, it is unclear which of these norm-belief pairs does the grounding work. And because these particular norm-belief pairs cannot be directly compared, I proposed that we search for a plausible second-order norm that somehow adjudicates between them. This second-order norm, together with whatever norm-belief pair it picks out, would ground the alternatives' subjective deontic statuses.

I believe that the stringency norm is the second-order norm that we seek. Cases of apparent incommensurability fall within the subclass of special cases in which the stringency norm is capable of operating: (1) the agent is uncertain between two beliefs (here, she is OI-EI uncertain); (2) a different rational norm would be operative if the agent held one belief than if she held the other; and (3) those two norms are not directly comparable (in this case, because OI is not amenable to numeric representation).

I said in Chapter 1 that I would remain neutral between theories about which norm is operative in OI cases: maximizing, optimizing, or optimizing + parity. I will now walk through the stringency norm's implications for each of these three OI-norm theories. Under any of the three theories, the presence of OI-EI uncertainty and the operation of the stringency norm change our understanding of how practical rationality works in putative OI cases. In the case of the optimizing norm, this change is dramatic. For the other two norms, the change is more limited. But in all three cases the stringency norm has important and surprising implications for the practical rationality of value incommensurability.

3.3 – The Stringency Norm's Effect on Putative OI Cases...

3.3.1 – ...Under a Maximizing Theory of OI

In Chapter 1, I laid out three possible theories about the rational norm that grounds alternatives' deontic statuses in OI cases. One theory is that alternatives' deontic statuses are

grounded in a norm of maximizing. An alternative maximizes if and only if it is *not worse than* each other alternative. In Chapter 1 there was no need to distinguish objective and subjective deontic statuses, because I took it for granted that the agent knew she faced an OI (rather than EI) case. But if my arguments in the first half of this chapter succeed, then this is rarely if ever the case. In virtually all putative OI cases, some degree of OI-EI uncertainty is present (both normatively and descriptively).

This OI-EI uncertainty has no effect on alternatives' objective deontic statuses (see above, § 3.1.1), but it does have the potential to affect alternatives' subjective deontic statuses. In a case of apparent incommensurability with OI-EI uncertainty, there are two candidate norm-belief pairs. Assuming (for the moment) a maximizing theory of OI, the first norm-belief pair consists of: (1) a rational norm of maximizing; and (2) a belief that the alternatives are ontologically incommensurable. The second norm-belief pair consists of: (1) Maximin; and (2) a belief that the alternatives are epistemically incommensurable. Because the agent is uncertain between the two beliefs, and because the two norms cannot be directly compared, the stringency norm governs which norm-belief pair(s) ground(s) alternatives' subjective deontic statuses. Under the stringency norm, each norm-belief pair grounds subjective deontic statuses if it is at least as stringent as the other norm-belief pair.

Here is how the stringency norm operates in a two-alternative $\{A,B\}$ apparent incommensurability case, under a maximizing theory of OI. Under the OI norm-belief pair, both A and B are subjectively rationally permissible, because both alternatives maximize value: given the belief that the alternatives are OI, neither is worse than the other (see Chapter 1). Under the EI norm-belief pair, both A and B are subjectively rationally permissible, because both alternatives satisfy Maximin: given the belief that the alternatives are EI, each maximizes

minimum possible value (see Chapter 2). A and B would thus have identical subjective deontic statuses under either norm-belief pair. Both norm-belief pairs are therefore equally stringent: neither renders subjectively rationally impermissible an alternative that the other renders subjectively rationally permissible. Under the stringency norm, both norm-belief pairs, together with the stringency norm, ground the alternatives' subjective deontic statuses. The end result is that both A and B are subjectively rationally permissible. The same is true in cases of complete apparent incommensurability with more than two alternatives: every alternative is subjectively rationally permissible.

This is a sensible and intuitive, if not particularly surprising, result. But stringency holds relative to a set of alternatives. The fact that two norm-belief pairs are equally stringent relative to one set of alternatives does not guarantee that they will be equally stringent relative to *every* set of alternatives. In cases of partial apparent incommensurability, the results are sometimes more surprising. Recall the National Book Award example from Chapter 2, in which A appears incommensurable with B and with C , but B appears slightly better than C . Under the OI norm-belief pair (assuming a maximizing theory of OI), both A and B are subjectively rationally permissible: given the belief that the alternatives are OI, neither A nor B is worse than any other alternative. (C is subjectively rationally impermissible, because it is worse than B .) But under the EI norm-belief pair, B is subjectively rationally obligatory: given the belief that the alternatives are EI, B maximizes minimum possible value (see Chapter 2).

Table 3.9

	A	B	C
EI norm-belief pair	impermissible	permissible	impermissible
OI norm-belief pair	permissible	permissible	impermissible

In this choice, the EI norm-belief pair is more stringent than the OI norm-belief pair.

Given two norm-belief pairs $\{x,y\}$, x is more stringent than y if and only if: (1) there is at least

one alternative that x , but not y , renders subjectively rationally impermissible; but (2) there is no alternative that y , but not x , renders subjectively rationally impermissible. Here, the EI norm-belief pair, but not the OI norm-belief pair, renders A subjectively rationally impermissible. On the other hand, there is no alternative that the OI norm-belief pair renders subjectively rationally impermissible but the EI norm-belief pair renders permissible. Because the EI norm-belief pair is more stringent than the OI pair, it grounds the alternatives' subjective deontic statuses (together with the stringency norm). The result is that the committee is subjectively rationally required to give the award to B .

This should be a surprising result for adherents of the maximizing theory of OI—particularly if this choice, or some other choice whose alternatives appear partially incommensurable in the same way,⁹⁶ strikes them as an intuitively plausible instance of OI. OI-EI uncertainty has the effect of rendering an alternative subjectively rationally impermissible, *even though it appears to maximize value*. This is true regardless of the degree of OI-EI uncertainty, because the stringency norm does not take the OI and EI beliefs' subjective probabilities into account (see above, § 3.2). So long as the agent has (or should have) a nonzero subjective probability both that she faces an OI case and that she faces an EI case, she ought obey the more stringent norm—even if she has only a *de minimis* credence that she faces an EI case. And since she will almost always have at least a *de minimis* credence that she faces an EI case, she ought almost always obey the more stringent norm.

To sum up, the effect of the stringency norm on the maximizing theory of OI is that in putative OI cases, alternatives will sometimes be subjectively rationally impermissible, even though they appear to maximize value. This is a surprising and important result, though it

⁹⁶ That is, there are three alternatives $\{A, B, C\}$ such that A appears incommensurable with either B or C , but B appears to be better than C .

admittedly has limited application. It holds only in some cases of partial apparent incommensurability, and does not hold in any cases of complete apparent incommensurability.

3.3.2 – ...*Under an Optimizing Theory of OI*

A second theory of OI is that alternatives' deontic statuses are grounded in an optimizing norm. An alternative optimizes if and only if it is *at least as good as* each other alternative. Under an optimizing theory of OI, OI-EI uncertainty and the stringency norm have a more dramatic effect on subjective deontic statuses.

Here is how the stringency norm operates in a two-alternative $\{A,B\}$ case of apparent incommensurability under an optimizing theory of OI. As before, the EI norm-belief pair renders both alternatives subjectively rationally permissible. The OI norm-belief pair, on the other hand, fails to render any deontic statuses at all. Under the optimizing norm, an alternative is permissible if and only if it is at least as good as each other alternative. When paired with the belief that neither alternative is at least as good as the other, the optimizing norm fails to render either alternative subjectively rationally permissible. But, as I discussed in Chapter 1, we would espouse an inaccurate characterization of the optimizing norm (at least as that norm has been portrayed by other philosophers) if we inferred from the fact that neither alternative is permissible that both alternatives are therefore impermissible. In cases of apparent incommensurability, proponents of the optimizing norm hold that reasons “run out.” Practical reason simply has no guidance to offer to the agent, and she is forced to choose by “plumping”—that is, arbitrarily selecting an alternative with no guidance from practical reason.⁹⁷ Hence, it is more accurate to state the optimizing norm's conditions for rational impermissibility as follows: an alternative is *impermissible* if and only if it is *worse than* each other alternative. If we

⁹⁷ For a helpful discussion of how “plumping” differs from the “picking” that agents do when they select one of several rationally permissible alternatives, see Chang 2017, § 3.

understand the optimizing norm in this way—as I argued in Chapter 1 that we should—then the optimizing norm fails to render any subjective deontic statuses when paired with the belief that *A* and *B* are ontologically incommensurable. It is silent.

The OI norm-belief pair renders no subjective deontic statuses, and the EI norm-belief pair renders both alternatives subjectively rationally permissible. Norm-belief pairs are equally stringent if and only if neither renders an alternative (subjectively rationally) impermissible that the other renders (subjectively rationally) permissible. Here, neither the OI norm-belief pair nor the EI norm-belief pair renders any alternatives subjectively rationally impermissible. This is an instance of equal stringency—albeit a less straightforward one than we encountered under the maximizing theory of OI—and so both norm-belief pairs jointly ground the alternatives’ subjective deontic statuses (together with the stringency norm).

Under the maximizing theory of OI, both norm-belief pairs rendered identical subjective deontic statuses. They were, in a sense, redundant—or at least in agreement. Here, the norm-belief pairs do not render identical subjective deontic statuses. But neither do they render logically inconsistent subjective deontic statuses. Rather, the OI norm-belief pair fails to render any subjective deontic status at all. Meanwhile, the EI norm-belief pair does render a set of subjective deontic statuses. Both norm-belief pairs may thus ground the subjective deontic statuses without yielding a contradiction. Speaking metaphorically, the EI norm-belief pair casts a vote in favor of one set of subjective deontic statuses, while the OI norm-belief pair abstains from voting. Since the stringency norm holds that their two votes jointly determine the alternatives’ subjective deontic statuses, the EI norm-belief pair’s vote carries the day. For the same reasons, the EI norm-belief pair’s vote will carry the day in cases of complete apparent incommensurability with more than two alternatives.

In a case of partial apparent incommensurability such as the National Book Award example, the EI norm-belief pair is more stringent than the OI norm-belief pair. According to the EI norm-belief pair, book *B* is subjectively rationally obligatory, while books *A* and *C* are subjectively rationally impermissible. The OI norm-belief pair, on the other hand, once again fails to render any subjective deontic statuses. No alternative is subjectively rationally permissible, because no alternative is at least as good as each other alternative; but no alternative is subjectively rationally impermissible, because no alternative is worse than each other alternative. Because the EI norm-belief pair (but not the OI norm-belief pair) renders *A* and *C* subjectively rationally impermissible, and because there is no alternative that only the OI norm-belief pair renders subjectively rationally impermissible, the EI norm-belief pair is more stringent than the OI norm-belief pair. Thus, the EI norm-belief pair grounds the alternatives' subjective deontic statuses (together with the stringency norm).

In either case—whether the EI and OI norm-belief pairs are equally stringent (as in cases of complete apparent incommensurability), or whether the EI norm-belief pair is more stringent than the OI norm-belief pair (as in cases of partial apparent incommensurability)—OI-EI uncertainty effectively ensures that apparently incommensurable alternatives have subjective deontic statuses, even under an optimizing theory of OI. OI-EI uncertainty triggers the stringency norm, which holds that the EI norm-belief pair grounds alternatives' subjective deontic statuses—either jointly with the OI norm-belief pair (if the norm-belief pairs are equally stringent) or by itself (if it is more stringent than the OI norm-belief pair). Unlike the OI norm-belief pair, the EI norm-belief pair is not silent: it renders subjective deontic statuses. And since OI-EI uncertainty is present in virtually all cases of apparent incommensurability, there will be virtually no choices in which the alternatives fail to have subjective deontic statuses.

I said in Chapter 1 that under the optimizing theory of OI, practical rationality breaks down in OI cases. Happily, the presence of OI-EI uncertainty and the operation of the stringency norm allow us to substantially temper that claim. While practical rationality breaks down at the level of objective rightness, it rarely (if ever) breaks down at the level of subjective rightness. This is welcome news, because from a practical standpoint, it is generally subjective rightness that matters. Even if our alternatives turn out to be ontologically incommensurable, we are not reduced to mere plumping.

3.3.3 – ...Under a Parity Theory of OI

The third theory of OI is that alternatives' deontic statuses are grounded in a version of the optimizing norm that has been expanded to include parity. Parity (or "on a par"), according to its proponents, is a fourth value relation that is mutually exclusive with "better than," "worse than," or "equally good." Because "on a par" is mutually exclusive with the three standard value relations, on-a-par alternatives are ontologically incommensurable. (This follows from the definition of OI—two alternatives are OI if and only if none of the three standard value relations obtains between them.) Parity's proponents endorse what I've dubbed the "optimizing + parity" norm: an alternative is rationally permissible (absent deontic constraints) if and only if it is *either* at least as good as *or on a par with* each other alternative.

Here is how the stringency norm operates in a two-alternative $\{A,B\}$ case of apparent incommensurability under the theory that ontologically incommensurable alternatives are on a par. Under the OI norm-belief pair, (1) A and B are on a par, and (2) according to the optimizing + parity norm, both are subjectively rationally permissible. Under the EI norm-belief pair, both are also subjectively rationally permissible (as previously discussed). The OI and EI norm-belief pairs are thus equally stringent; they both ground the alternatives' subjective deontic statuses,

together with the stringency norm. The same is true of cases of complete apparent incommensurability that involve more than two alternatives.

In a case of partial apparent incommensurability such as the National Book Award example, the EI norm-belief pair is more stringent than the OI norm-belief pair. Under a parity interpretation of the example, *A* is on a par with *B* and with *C*, but *B* is slightly better than *C*. *A* is subjectively rationally permissible, because it is on a par with each other alternative. *B* is subjectively rationally permissible, because it is on a par with *A* and better than *C*. *C* is subjectively rationally impermissible, because it is worse than *B*. But under the EI norm-belief pair, *A* and *C* are both subjectively rationally impermissible—only *B* is subjectively rationally permissible.

Table 3.10: partial apparent incommensurability under a parity theory of OI

	<i>A</i>	<i>B</i>	<i>C</i>
EI norm-belief pair	impermissible	permissible	impermissible
OI norm-belief pair	permissible	permissible	impermissible

The EI norm-belief pair is thus more stringent than the OI norm-belief pair, because (1) there is an alternative—*A*—that the EI, but not the OI, norm-belief pair renders subjectively rationally impermissible, but (2) there is no alternative that the OI, but not the EI, norm-belief pair renders subjectively rationally impermissible. Because the EI norm-belief pair is more stringent, it alone grounds the alternatives' subjective deontic statuses (together with the stringency norm), and the committee is subjectively rationally required to give the award to book *B*.

The surprising result here is that an alternative may sometimes be subjectively rationally impermissible *even though it is (objectively speaking) on a par with each other alternative*—so long as there is some degree of OI-EI uncertainty. (And, as I argued earlier, there will almost always be some degree of OI-EI uncertainty.) Although this is surprising, I also think it makes good intuitive sense. The agent is uncertain whether she faces a parity case or an EI case. If it's

an EI case, the norm-belief pair that governs her choice is stricter—specifically, more stringent—than the norm-belief pair that would govern if it were a parity case. Since she is unsure which norm-belief pair governs, she ought follow the stricter, more stringent pair—particularly since, by following the stricter pair, she will simultaneously obey the laxer, less stringent pair’s prescriptions. And because she *ought* to follow the stricter norm-belief pair—in other words, is subjectively rationally required to do so—it is in fact true that the stricter pair governs, in the sense that it grounds her alternatives’ subjective deontic statuses.

3.4 – Conclusion

This is a dissertation about practical rationality in cases of apparent incommensurability: those vexing choices in which our alternatives continue to appear incommensurable, no matter how thoroughly and competently we have deliberated, and despite a good grasp of all relevant empirical facts. In the preceding pages, I have offered my answer to the following question: when forced to choose between apparently incommensurable alternatives, what should we do, and why?

Cases of apparent incommensurability are either OI cases or EI cases. OI and EI are mutually exclusive and exhaustive. In Chapter 1, I explained the three leading theories of what grounds alternatives’ deontic statuses in OI cases. One theory is that alternatives’ deontic statuses are grounded in a maximizing norm, according to which an alternative is rationally permissible if and only if it is not worse than each other alternative. Because ontological incommensurability between alternatives entails that they are not worse than one another, every member of a set of ontologically incommensurable alternatives is rationally permissible. The second theory holds that maximizing is not sufficient to rationally justify choice. Alternatives’ deontic statuses are instead grounded in an optimizing norm, according to which an alternative is

rationally permissible if and only if it is at least as good as each other alternative. Unfortunately, the optimizing theory of OI entails that if an agent faces a set of ontologically incommensurable alternatives, no alternative is rationally permissible. More recently, several philosophers have argued that we can avoid this unfortunate result not by accepting a maximizing theory of OI, but instead by positing a new value relation, parity, that can hold between ontologically incommensurable alternatives. Parity's defenders endorse an optimizing + parity norm, under which an alternative is rationally permissible if and only if it either optimizes or is on a par with each other alternative. Since OI has already been well-examined by others, I stated that I would remain neutral between these three theories and focus instead on aspects of value incommensurability that the existing literature has neglected.

In Chapter 2, I offered an account of what grounds alternatives' deontic statuses in EI cases. I distinguished between objective and subjective deontic statuses, and argued that while the MAV norm—a norm of maximizing actual value—plausibly grounds the former, it does not ground the latter. I then examined at some length whether the MEV norm—a norm of maximizing expected value—might instead ground alternatives' subjective deontic statuses. I concluded that it cannot do so, even with the help of the principle of insufficient reason. I went on to argue that alternatives' subjective deontic statuses in EI cases are instead grounded in Maximin—a norm of maximizing minimum possible value.

In this chapter (Chapter 3), I turned my attention to the question of how OI-EI uncertainty affects subjective deontic statuses in cases of apparent incommensurability. (OI-EI uncertainty does not affect objective deontic statuses, because objective rightness is not a function of the agent's beliefs.) I argued that OI-EI uncertainty figures into virtually all cases of apparent incommensurability. Since I wanted to remain neutral on the question of whether subjective

rightness is a function of actual or reasonable beliefs, I argued that OI-EI uncertainty is both descriptively present (i.e., it is an actual belief) and normatively required (i.e., it is a reasonable belief) in cases of apparent incommensurability. I then proposed that OI-EI uncertainty affects subjective deontic statuses via the stringency norm.

I finished by examining how the stringency norm affects subjective deontic statuses in cases of apparent incommensurability under each of the three theories of OI from Chapter 1. My goal in doing so was to provide an account of practical rationality under OI-EI uncertainty whose usefulness is not tied to any particular OI-norm theory. On the other hand, I did tie my account to a particular theory of which rational norm governs EI cases. This is because little has been said elsewhere about the practical rationality of EI cases. Because there aren't any extant competing theories, I was forced to rely solely on my own account of EI cases from Chapter 2.

The stringency norm's effects are most dramatic under an optimizing theory of OI. Because the optimizing norm-belief pair fails to render any deontic statuses, the EI norm-belief pair will always be at least as stringent as the OI norm-belief pair. Hence the EI norm-belief pair grounds alternatives' subjective deontic statuses, either jointly with the silent OI norm-belief pair (when the pairs are equally stringent) or singly (when the EI norm-belief pair is more stringent than the OI norm-belief pair). Even under an optimizing theory of OI, then, alternatives will have subjective deontic statuses, so long as there is some degree of OI-EI uncertainty. And since some degree of OI-EI uncertainty is virtually always present, alternatives will virtually always have subjective deontic statuses. Practical rationality does not break down—at least at the level of subjective rightness, which is the level that is relevant to guiding agents' decisions. Agents are not reduced to mere plumping.

The presence of OI-EI uncertainty and the operation of the stringency norm also sometimes affect subjective deontic statuses under the maximizing and parity theories of OI. In cases of complete apparent incommensurability, the stringency norm does not affect the alternatives' subjective deontic statuses. Because the EI norm-belief pair and the OI norm-belief pair (whether under a maximizing or optimizing + parity norm) both render identical subjective deontic statuses (specifically, they render every alternative subjectively rationally permissible), the alternatives' subjective deontic statuses will be the same under either. In cases of *partial* apparent incommensurability, on the other hand, the effect of OI-EI uncertainty (via the stringency norm) is that an alternative may sometimes be subjectively rationally impermissible despite the fact that it (1) maximizes (under a maximizing theory of OI) or (2) is on a par with each other alternative (under a parity theory of OI).

An interesting result is that in cases of both complete and partial apparent incommensurability, under any of the three theories of OI, the EI norm-belief pair is always at least as stringent as the OI norm-belief pair. Thus, as a heuristic matter, an agent can simply proceed as if she faces an EI case (even if she in fact faces an OI case).

Together, the three chapters comprise a complete account of what grounds alternatives' deontic statuses in cases of apparent incommensurability (absent deontic constraints, time constraints, empirical ignorance, and an incompetent or disingenuous practical reasoner). The account covers each half of a mutually exclusive and exhaustive subdivision of cases of apparent incommensurability into OI and EI cases, and also covers how practical reason operates when there is uncertainty as to which half of the subdivision applies in the choice at hand.

Bibliography

Anderson, Elizabeth, 1993, *Value in Ethics and Economics*, Cambridge: Harvard University Press.

Anderson, Jack, 2015, “Resolving the Small Improvement Argument: A Defense of the Axiom of Completeness,” *Erasmus Journal for Philosophy and Economics*, 8 (1): 24-41.

Andersson, Henrik, “Parity and Comparability—A Concern Regarding Chang’s Chaining Argument,” *Ethical Theory and Moral Practice*, 19 (1): 245-253.

Bliss, Ricki and Trogdon, Kelly, “Metaphysical Grounding,” *The Stanford Encycloedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/win2016/entries/grounding/>](https://plato.stanford.edu/archives/win2016/entries/grounding/).

Bradley, Seamus, “Imprecise Probabilities,” *The Stanford Encyclopedia of Philosophy* (Winter 2016 edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/win2016/entries/imprecise-probabilities/>](https://plato.stanford.edu/archives/win2016/entries/imprecise-probabilities/).

Briggs, R.A., 2017, “Normative Theories of Rational Choice: Expected Utility,” *The Stanford Encyclopedia of Philosophy* (Spring 2017 edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/spr2017/entries/rationality-normative-utility/>](https://plato.stanford.edu/archives/spr2017/entries/rationality-normative-utility/).

Broome, John, 1997, “Is Incommensurability Vagueness?,” in *Incommensurability, Incomparability, and Practical Reason*, R. Chang (ed.), Cambridge: Harvard University Press.

Buchak, Lara, 2013, *Risk and Rationality*, Oxford: Oxford University Press.

Boot, Martijn, 2009, “Parity, Incomparability and Rationally Justified Choice,” *Philosophical Studies*, 146 (1): 75-92.

Chang, Ruth, (ed.), 1997a, *Incommensurability, Incomparability, and Practical Reason*, Cambridge: Harvard University Press.

Chang, Ruth, 1997b, “Introduction,” in *Incommensurability, Incomparability, and Practical Reason*, R. Chang (ed.), Cambridge: Harvard University Press.

Chang, Ruth, 1998, “Comparison and the Justification of Choice,” *University of Pennsylvania Law Review*, 146 (5): 1569-1598.

Chang, Ruth, 2002a, *Making Comparisons Count*, New York: Routledge.

Chang, Ruth, 2002b, “The Possibility of Parity,” *Ethics*, 112: 659-688.

Chang, Ruth, 2004a, “All Things Considered,” *Philosophical Perspectives*, 18 (1): 1-22.

Chang, Ruth, 2004b, "Putting Together Morality and Well-Being," in *Practical Conflicts*, M. Betzler and P. Baumann (eds.), Cambridge: Cambridge University Press.

Chang, Ruth, 2005, "Parity, Interval Value, and Choice," *Ethics*, 115: 331-350.

Chang, Ruth, 2012, "Are Hard Choices Cases of Incomparability?," *Philosophical Issues*, 22: 106-126.

Chang, Ruth, 2016a, "Comparativism: The Grounds of Rational Choice," in *Weighing Reasons*, E. Lord and B. Maguire (eds.), New York: Oxford University Press.

Chang, Ruth, 2016b, "Parity, Imprecise Comparability, and the Repugnant Conclusion," *Theoria*, 82: 182-214.

Chang, Ruth, 2017, "Hard Choices," *Journal of the American Philosophical Association*, 3 (1): 1-21.

Dubs, Homer, 1942, "The Principle of Insufficient Reason," *Philosophy of Science*, 9 (2): 123-131.

Editors of Encyclopædia Britannica, "Mohs Hardness," *Encyclopædia Britannica*, URL = <<https://www.britannica.com/science/Mohs-hardness>>.

Elson, Luke, 2014, "Heaps and Chains: Is the Chaining Argument for Parity a Sorites?," *Ethics*, 124: 557-571.

Elson, Luke, 2017, "Incommensurability as Vagueness: A Burden-Shifting Argument," *Theoria* 83: 341-363.

Espinoza, Nicolas, 2008, "The Small Improvement Argument," *Synthese*, 65: 127-139.

Gert, Joshua, 2004, "Value and Parity," *Ethics*, 114: 492-510.

Gert, Joshua, 2015, "Parity, Preference, and Puzzlement," *Theoria*, 81: 249-271.

Griffin, James, 1986, *Well-Being: Its Meaning, Measurement, and Importance*, Oxford: Clarendon Press.

Gustaffson, Johan E., 2013, "Indeterminacy and the Small-Improvement Argument," *Utilitas* 25 (4): 433-445.

Gustaffson, Johan E. and Nicolas Espinoza, 2010, "Conflicting Reasons in the Small-Improvement Argument," *The Philosophical Quarterly*, 60 (241): 754-763.

Harman, Elizabeth, 2015, "The Irrelevance of Moral Uncertainty," R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 10*, New York: Oxford University Press, pp. 53-79.

Hsieh, Nien-hê, 2005, "Equality, Clumpiness and Incomparability," *Utilitas*, 17 (2): 180-204.

Hsieh, Nien-hê, 2007, "Is Incomparability a Problem for Anyone?," *Economics and Philosophy*, 23: 65-80.

Huber, Franz, 2009, "Belief and Degrees of Belief," in *Degrees of Belief*, F. Huber and C. Schmidt-Petri (eds.), Dordrecht: Springer.

Jeffrey, Richard, 1983, *The Logic of Decision*, 2nd edition, Chicago: University of Chicago Press.

Kelly, Chris, 2008, "The Impossibility of Incommensurable Values," *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 137 (3): 369-382.

Klocksiem, Justin, 2010, "In Defense of the Trichotomy Thesis," *Acta Analytica*, 25 (3): 317-327.

Kruger, Justin and David Dunning, 1999, "Unskilled and Unaware of It: How Difficulties in Recognizing One's Own Incompetence Lead to Inflated Self-Assessments," *Journal of Personality and Social Psychology*, 77 (6): 1121-1134.

Lockhart, Ted, 2000, *Moral Uncertainty and Its Consequences*, New York: Oxford University Press.

Luce, R.D., and H. Raiffa, 1957, *Games and Decisions*, New York: Wiley.

Mendola, Joseph, 1990, "An Ordinal Modification of Classical Utilitarianism," *Erkenntnis*, 33 (1): 73-88.

Mendola, Joseph, 2005, "Intuitive Maximin," *Canadian Journal of Philosophy*, 35 (3): 429-439.

Mendola, Joseph, 2006, *Goodness and Justice*, New York: Cambridge University Press.

Mendola, Joseph, 2014, "The Leximin Desire Principle," Chapter 8 in *Human Interests: or Ethics for Physicalists*, New York: Oxford University Press.

Moon, David, unpublished a, "Rough Equality and Comparability: Problems With Existing Proposals."

Moon, David, unpublished b, "Value Incomparability and Practical Reason."

Moore, G.E., 1903, *Principia Ethica*, Cambridge: Cambridge University Press.

Nozick, Robert, 1974, *Anarchy, State, and Utopia*, New York: Basic Books.

Peterson, Martin, 2009, *An Introduction to Decision Theory*, Cambridge: Cambridge University Press.

Rabinowicz, Wlodek, 2008, "Value Relations," *Theoria*, 74 (1): 18-49.

Rabinowicz, Wlodek, 2009a, "Incommensurability and Vagueness," *Aristotelian Society Supplementary*, 83 (1): 71-94.

Rabinowicz, Wlodek, 2009b, "Values Compared," *Polish Journal of Philosophy*, 3 (1): 73-96.

Rabinowicz, Wlodek, 2012, "Value Relations Revisited," *Economics and Philosophy*, 28 (2): 133-164.

Rapoport, Anatol, 1989, *Decision Theory and Decision Behavior*, Dordrecht: Springer.

Rawls, John, 1971, *A Theory of Justice*, Cambridge, MA: Harvard University Press. Revised edition, 1999.

Raz, Joseph, 1986, *The Morality of Freedom*, Oxford: Clarendon Press.

Raz, Joseph, 1997, "Incommensurability and Agency," in R. Chang (ed.), *Incommensurability, Incomparability, and Practical Reason*, Cambridge: Harvard University Press, pp. 110-128.

Raz, Joseph, 1999, *Practical Reason and Norms*, Oxford: Oxford University Press.

Regan, Donald, 1997, "Value, Comparability, and Choice," in R. Chang (ed.), *Incommensurability, Incomparability, and Practical Reason*, Cambridge: Harvard University Press, pp. 129-150.

Resnik, Michael, 1987, *Choices: An Introduction to Decision Theory*, Minneapolis: University of Minnesota Press.

Ross, Jacob, 2006, "Rejecting Ethical Deflationism," *Ethics*, 116 (4): 742-768.

Savage, Leonard J., 1954, *The Foundations of Statistics*, New York: Wiley.

Schotter, Andrew, 2008, *Microeconomics: A Modern Approach*, Mason: Cengage Learning.

Sen, Amartya, 1997, "Maximization and the Act of Choice," *Econometrica*, 65: 745-779.

Sen, Amartya, 2000, "Consequential Evaluation and Practical Reason," *Journal of Philosophy*, 98: 477-502.

Sepielli, Andrew, 2009, "What To Do When You Don't Know What To Do," in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 4*, New York: Oxford University Press, pp. 5-28.

- Sepielli, Andrew, 2013, "Moral Uncertainty and the Principle of Equity Among Moral Theories," *Philosophy and Phenomenological Research*, 86 (3): 580-589.
- Sepielli, Andrew, 2014, "What To Do When You Don't Know What To Do When You Don't Know What To Do ...," *Noûs*, 48(3): 521-544.
- Simon, Herbert, 1955, "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics*, 69: 99-118.
- Sinn, Hans-Werner, 1980, "A Rehabilitation of the Principle of Insufficient Reason," *The Quarterly Journal of Economics*, 94 (3): 493-506.
- Slote, Michael, 1989, *Beyond Optimizing: A Study of Rational Choice*, Cambridge: Harvard University Press.
- Smith, Holly, 2010a, "The Moral Clout of Reasonable Beliefs," in M. Timmons (ed.), *Oxford Studies in Normative Ethics, Vol. I*, Oxford: Oxford University Press.
- Smith, Holly, 2010b, "Subjective Rightness," *Social Philosophy and Policy*, 27 (2): 64-110.
- Temkin, Larry, 2012, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*, Oxford: Oxford University Press.
- van Fraassen, Bas, 1989, *Laws and Symmetry*, Oxford: Oxford University Press.
- von Neumann, J. and O. Morgenstern, 1944, *Theory of Games and Economic Behavior*, Princeton: Princeton University Press.
- Waldron, Jeremy, 2003, "Security and Liberty: The Image of Balance," *The Journal of Political Philosophy*, 11: 191-210.
- Wasserman, Ryan, 2004, "Indeterminacy, Ignorance and the Possibility of Parity," *Philosophical Perspectives*, 18: 391-403.