Formal Epistemology without Demandingness*

Tim Smartt
timjsmartt@gmail.com

Forthcoming in *Synthese*

Abstract

I argue that the methodology of model building motivates the view that the norms of formal epistemology should not be excessively demanding. This is quite a different picture than one often encounters, especially among philosophers who are sceptical of the usefulness of formal work in epistemology. I argue for this view in two ways. First, formal epistemologists are engaged in a particular kind of modelling—namely, normative modelling—which includes a feature that *supports* demandingness objections. One role of normative models is to deliver normative guidance—indeed, this is a key difference between normative models and descriptive models. Of course, *should* and *can* are two different matters, and one might wonder: even if normative models should be action-guiding, *can* they be? Second, I argue that normative epistemological models can deliver on this promise. Taken together, the lines of argument make the case that normative models *should* and *can* deliver normative guidance. In other words, we can have formal epistemology without demandingness.

Keywords: Formal epistemology · Modelling · Normative modelling

*Thanks to Brian Hedden, Mark Colyvan, and two anonymous referees for very helpful comments that improved this paper.

1 Introduction

Can epistemology require the impossible, or even the very difficult? For instance, can it require us to believe all and only the truths, or to have only point-valued credences uniquely picked out by the evidence, or to get ourselves into a coherent total doxastic state? One reason why you might think the answer is 'no' is because you take it that there is a demandingness constraint on epistemic norms. Sure, epistemology should be demanding in some sense—no doubt we are required to do some things to responsibly manage our doxastic lives—but there are limits. Epistemic requirements that surpass these limits are excessively demanding. These kinds of demandingness objections are common in recent epistemology. But so too are advocates of highly demanding norms. How might they respond?

One way to defend demanding norms in epistemology goes like this—I will call it *the modelling defence*. Epistemology can be done in a formal or a traditional register. Although demandingness objections are directed towards norms in both traditions, they tend to target the norms of formal epistemology—norms of rationality like probabilism, or norms of belief revision like conditionalization. But formal epistemology employs the distinctive methodology of model building, and demandingness objections that target formal epistemology misunderstand this methodology. Formal models include elements that can look daunting, such as ideal agents, abstract objects, and mathematical formulae. Demandingness objections directed at formal epistemology latch onto these elements and express worries about a picture of human reasoning that seems far removed from the mental lives of ordinary thinkers. But—*the modelling defence* claims—these worries arise from misunderstanding the methodology of model building. For instance, models are not laws, often their abstract objects are not to be interpreted literally, and some of the idealizations and mathematics is just for convenience. When the methodology is better understood, these sorts of misunderstandings are overcome, and demandingness objections lose their force.

I argue that *the modelling defence* fails to provide a satisfying reply to demandingness objections directed at the norms of formal epistemology. Demandingness objections stick even when one better understands model building. This is not due to a fault with modelling—I am sympathetic to model building in epistemology, and in other areas of philosophy. It is because the methodology of model building *vindicates* demandingness objections rather

than overcoming them. Formal epistemologists are engaged in a particular kind of modelling—namely, normative modelling—which includes a feature that *supports* demandingness objections. Roughly, the feature is that one purpose of normative models is to deliver useful advice. So, formal epistemologists cannot avoid demandingness objections by appealing to the methodology of normative modelling. I argue that there is a different lesson to take from *the modelling defence*: the methodology of normative modelling motivates the view that the norms of formal epistemology should not be excessively demanding.

This paper is in three sections. Section 2 focuses on modelling. I clarify several methodological features of modelling, focusing especially on *normative* modelling as a subspecies of this approach. In the rest of the paper, I draw out an implication from the nature of normative modelling for the demandingness of the norms of formal epistemology. In Section 3, I show that distinguishing *normative* from *descriptive* modelling involves connecting normative modelling with some sort of 'ought.' In Section 4, I discuss a framework due to Titelbaum (2012) which can be used to appreciate how normative models deliver the relevant 'oughts.' Taken together, these sections present a picture according to which normative models both *can* and *should* deliver some sort of normative guidance. In both Section 3 and Section 4, I also respond to several objections that one might have to either claim; that is, to either the claim that providing advice is among the central purposes of normative models (Section 3) or to the claim that they can do so in an apt manner (Section 4). In a nutshell, this paper makes the case that formal epistemologists will not find a compelling response to demandingness objections from the nature of normative modelling. A central feature of normative modelling supports the idea that epistemic norms should not be excessively demanding.

My view has several metaepistemological upshots for broader debates about formal versus traditional epistemology. For starters, it reveals one way in which there is not a clean distinction between the two approaches. It is often assumed that the demand for usable normative guidance comes from traditional epistemology. But I show that methodological features of formal epistemology support these kinds of considerations too, so the notion of guidance is not a genuine fault line. More broadly, my view suggests a way in which standard ways of thinking about the difference between formal and traditional epistemology can be improved. My view outlines a novel way of combining metaepistemological intuitions that usually lead in different directions. Standardly, if you think that epistemology should not be too demanding, you

will opt for traditional epistemology, perhaps in an especially 'non-ideal' vein. Likewise, if you are sympathetic to model building, you will opt for formal epistemology. I am sympathetic to both intuitions and think that they can be embraced at the same time. In a slogan, this paper proposes that we can have formal epistemology without demandingness.

It is worth emphasising that the goal of this paper is *not* to present a criticism of formal epistemology or normative modelling. Just the opposite: I hope to suggest new ways of advocating for these approaches to philosophers who have doubts about how useful they are. On the picture that I explore, the methodology of formal modelling naturally puts constraints on how demanding the norms of formal epistemology ought to be. This paper only represents a challenge to formal epistemology or normative modelling if one takes it that demanding norms are an essential feature of these approaches. And I hope to show that this is a commitment that we have good reason to resist.

2 On normative modelling

What is normative modelling? It is best understood by distinguishing it from two other theoretical approaches; namely, theory building and descriptive modelling. Let's take these in turn.

Modelling, in general, is a particular approach to studying a phenomenon. It is common right across the natural and social sciences—from economics and political science to biology and physics. It has many of the same general aspirations as other forms of inquiry. For instance, it aims at capturing some aspect of reality and allowing us to better understand it. But modelling is different in a number of respects to what we might call 'theory building,' which tends to emphasise a way of capturing some aspect of reality that is uniquely true, fairly general and complete, and comprised of a set of axiomatizable principles (e.g., Williamson 2017, Roussos 2022, 2025).

Two significant points of difference are that models are *partial* and *indirect*. Theory building hopes to discover universal laws of nature, non-accidentally exceptionalness generalisations which can offer a basic and unifying explanation of the target phenomenon (Williamson 2017, 1).[1] Theories aim at being

[1]Are laws of nature really exceptionless? This question touches on debates in the philosophy of science about *ceteris paribus* laws; that is, purported scientific laws that only hold in certain circumstances (e.g., Earman and Roberts 1999, Earman, Roberts, and Smith 2002). For my purposes, we just need to note the fact that theoreticians hope to uncover universal laws, whereas model-

true, general, and complete. Models are much more targeted and restricted; they provide something more like a partial map, which illuminates some aspect of the target being studied. They aim at being fit to play a carefully specified role (Colyvan 2013, Appiah 2017). They do this by offering an *indirect* representation of the target phenomenon. Model-builders describe a model system and take it as the focus of their inquiry (Weisberg 2007, Godfrey-Smith 2007). When there is a good fit between a model and its target, exploring the model is an indirect way of exploring the target (Williamson 2017). This is especially useful when the target—like many aspects of the human and natural world—is sufficiently large and complex that it resists being captured by a theory. Models tend to involve mathematical formulae, simple and abstract environments—such as perfectly competitive markets (Rodrik 2015, 13), complete compliance (Ismael 2016), or predators who only eat one kind of prey (Colyvan 2014)—and fictional entities like ideal gases or perfectly rational and self-interested agents. In this sense, modelling aims at better understanding some aspect of a phenomenon by studying features that are known to not actually be a part of it (Godfrey-Smith 2009).

This is quite different to theory building. Modelling does not aim at producing the one and only universal model of a target phenomenon (Parker 2009), does not lead to the discovery of universal laws (Williamson 2017), does not claim to represent a target realistically or directly (Godfrey-Smith 2009), and delivers insights that need to be combined with those generated by a library of other models to gain an adequate picture of the target (Rodrik 2015, 65).

Normative modelling can be distinguished from descriptive modelling. While much of the methodology is the same, normative models are different in a few important respects. Most generally, normative models are those that are used in normative domains; domains where the primary theoretical issues are not understanding how, or explaining why, some system operates the way it does, but are rather about how it *ought* to operate.

A few writers have attempted to make the distinction between normative and descriptive modelling more precise. Colyvan (2013, 1338) holds that normative models have different aims to the explanatory or descriptive models that one tends to find in the natural and social sciences. Explanatory models allow us to understand how a target system works, descriptive models capture an aspect of the actual behaviour of a system, but normative models capture

builders do not. But it worth bearing in mind that how *universal* scientific laws are is an area of disagreement amongst philosophers of science.

how agents ought to behave. In the case of epistemology, normative models prescribe how agents ought to reason and organise their beliefs (Colyvan 2013, 1340). Similarly, Titelbaum (2025, 42-44) argues that the key difference between normative models and other sorts of models is the kind of systems that they target. Normative models attempt to fit normative facts such as prescriptions (e.g., 'you *should not* act against your own best interests'), evaluations (e.g., 'it is *irrational* to believe a contradiction'), and more general facts involving normative concepts (e.g., '*correct inference* requires truth-preservation'). On his view, formal epistemology, decision theory and formal ethics are paradigmatic examples of normative modelling. But he also suggests that some work in the philosophy of language and linguistics can be read as normative modelling, especially work aimed at capturing how rational speakers ought to communicate (e.g., Stalnaker 1978/2002).[2] Roussos (2025, 431) argues that the purposes of the modeller determine whether a model is normative or descriptive. Any model put to a normative use—say, evaluation or action-guidance—is a normative model. Roussos' view differs from Colyvan's and Titelbaum's since he takes it that normative models are not a fundamentally different kind to explanatory or descriptive models—sometimes the same model might be put to an explanatory use, sometimes a descriptive use, or sometimes a normative use. For my purposes, it will not matter which view we prefer from a range of more fine-grained ways of marking the distinction between normative and descriptive models. But in Section 3, I will return to these views to emphasise a significant point that they all *agree* on.

What is the difference between normative *modelling* and normative *theorising*? A philosopher sceptical of the theoretical utility of models in general might argue that normative models are especially easy to dispense with, since normative theories have all the same positive features that advocates of normative modelling claim for their own approach.[3] They might argue that normative theories identify how agents ought to behave, they capture and explain normative facts and concepts, and can be put to paradigmatically normative

[2] Although he does not discuss the distinction between normative and descriptive modelling, Williamson (2017, 6-8) also argues that several developments in the philosophy of language and linguistics are best understood as modelling rather than theorising.

[3] For example, Rosanna Keefe (2000, 49-61) makes an argument along these lines. She argues that model-based approaches to vagueness in the philosophy of language fail to answer the key questions that we want addressed on this topic (2000, 53). Furthermore, she claims that model-builders usually provide *ad hoc* accounts of which part of their models to take seriously, and they tend to use the claim *'It's only a model'* as an all-purpose escape clause to avoid problematic commitments (2000, 55). She advocates a theory building approach, claiming that a 'genuine theory' can deliver a true description of vague language (2000, 52).

purposes such as action-guidance and evaluation—they can do everything that Colyvan, Titelbaum and Roussos identify as the hallmarks of normative models. If this objection is right, then normative modelling looks like an unnecessary complication to introduce into one's philosophical methodology.

There are two questions here that it is worth disentangling. The first is whether the distinction between *modelling* and *theory building* holds up in the normative domain. The distinction was coined with scientific domains in mind, so it makes sense to consider whether it can be plausibly extended to normative domains like epistemology and ethics. The second question is whether normative modelling is preferable to normative theory building.

The second question seems to misunderstand what writers sympathetic to normative modelling tend to recommend. I am not aware of any philosopher who advocates *replacing* traditional philosophical theory building with modelling—for instance, Williamson (2017, 11) emphasises that model building is not the best or only way of doing philosophy, and Titelbaum (2025, 61) emphasises that normative modelling can be used in tandem with traditional theory building to pursue broader shared goals. Rather, advocates of normative modelling tend to recommend incorporating the methodology into philosophical work on certain topics where it might be useful for some purpose. Consider Jaakko Hintikka's work on epistemic logic. Hintikka (1962) develops models of belief and knowledge by adapting the formal apparatus of possible world semantics. But his models are not complete accounts of these epistemic states; they highlight particular aspects of the structure and dynamics of knowledge and belief that might complement a variety of other projects in epistemology. And, indeed, this is how they are usually received. For example, Williamson (2013) draws on Hintikka's models to examine the informal notion of 'justified true belief' at work in traditional Gettier cases. He develops model-based arguments that support the original intuition-based arguments due to Gettier (1963). Williamson (2013, 13) emphasises the fruitfulness of having methodologically diverse arguments converge on the same result, writing, 'One can arrive at the same conclusion either way, and each method lends support to the other. The Gettier effect is robust.'

We can answer the first question—that is, whether the distinction between modelling and theory-building holds up in the normative domain—by drawing on the kind of features that distinguish modelling from theory building in general. These features illustrate significant ways in which a model of a normative phenomenon differs from a theory of it. One such feature is *incom-*

pleteness—theories aim at capturing their target fully (at some level of generality), whereas models do not. Consider a couple of paradigmatic examples of normative theories: theories of right action and theories of linguistic meaning. These aim at providing a general and comprehensive account of moral rightness and meaningful expressions. One consequence of this aim is that theories must treat counterexamples with a degree of methodological seriousness. If we can provide a thought experiment in which an action seems morally right, or some speech seems meaningful, but the theory fails to deliver this verdict, advocates of the theory must offer a compelling response or else we have a good reason to think that the theory is false. But normative models do not have the same goals—they do not aim at uncovering general principles for some normative domain, which is why advocates of modelling emphasise that counterexamples play a much smaller role in this approach (Williamson 2017, 9; Titelbaum 2025, 55). A model of right action or linguistic meaning illuminates its target in a more partial manner.[4]

So far, I have made a four-way distinction between two kinds of theories and two kinds of models. An example might make this clearer. Take standard decision theory (i.e., expected utility theory). Here is what it would look like to interpret expected utility theory as *(i)* a descriptive theory, *(ii)* a normative theory, (*iii)* a descriptive model, and *(iv)* a normative model. This example also serves to highlight that choosing between these interpretations makes a difference, since standard decision theory is most plausible when read as a normative model.

We can take the two kinds of theories at once. Understood as either a *descriptive* or *normative theory* of rational choice under conditions of uncertainty, standard decision theory is an inadequate theory. Theories are supposed to be general and complete, but when one considers actual choice scenarios under uncertainty there are lots of ways that standard decision theory is incomplete. For instance, standard decision theory leaves out many familiar kinds of uncertainty. Sometimes we are not only uncertain what state of the world will

[4]I want to be careful not to overstate this point. Counterexamples still play a role in model building, especially when the counterexamples address a model's intended use. Consider the canonical counterexamples to expected utility theory, such as the St. Petersburg, Allais, and Ellsberg paradoxes. These are often taken as counterexamples to expected utility theory, since they reveal a clash between the model's verdict of what is the rational thing to do, and what we have strong independent grounds to think is the rational thing to do. My point here—which is also emphasised in Williamson (2017, 9) and Titelbaum (2025, 55)—is just that counterexamples play a smaller role in modelling, and can often be responded to differently, than in traditional philosophical theory building.

obtain, but also about other elements that are required to apply standard decision theory. For instance, sometimes we are uncertain about how best to frame a decision problem (Broome 1991, 95-107). Similarly, some kinds of uncertainty resist being measured in probabilistic terms, such as uncertainty arising from vagueness (Colyvan 2008). Read as a descriptive or normative theory, standard decision theory is also inadequate because it contains false elements, whereas theories are supposed to be true. Standard decision theory describes agents who are maximally opinionated (they have priors for all propositions in the algebra, and determinate credences and preferences about all the relevant prospects), logically omniscient, risk-neutral, and perfectly rational. Actual agents like you and me do not have these properties, so taken as a descriptive theory of rational choice under uncertainty—where its elements are to be read literally and evaluated as true or false—it is inadequate. And taken as a normative theory, it brings with it some other implausible implications: it evaluates all actual agents as irrational *and* it leaves out a range of relevant issues that we can reasonably expect a theory to address.[5]

Taken as a *descriptive model* of rational choice under uncertainty, standard decision theory is to be understood as a way of capturing important features of how actual agents in these conditions are disposed to make decisions, or as an explanation of the mechanisms at work in their decision-making, or as a way of predicting what they will choose in a particular scenario. Standard decision theory is usually not apt for these kinds of purposes. It often does not model the psychological mechanisms involved in decision making, or predict how actual agents will make decisions, with much accuracy. Theorists interested in these types of questions have developed alternative models that are apt for descriptive, explanatory, and predictive purposes (e.g., Kahneman and Tversky 1979, Tversky and Kahneman 1992). It is as a *normative* model that standard decision theory is usually interpreted and defended; that is, as a model that captures a cluster of core ideas regarding how agents ought to make decisions if they are to be rational. This can be distinguished from a normative *theory* since it contains features that are hallmarks of models, such as elements that are false if interpreted literally, a partial view of its target phenomenon, and

[5]It bears emphasising that the only difference between interpreting standard decision theory as a descriptive theory or a normative theory is whether one takes it that it captures how people actually make decisions or how they ought to (if they are to be rational). Although the second interpretation is more familiar, it is just as inadequate as the first when understood as an instance of theory building, and for largely the same reasons. Normative theories that contain false elements, or provide an incomplete account of the target phenomenon, are just as flawed as descriptive theories that have the same failings.

limited domains of applicability. And it can be distinguished from a *descriptive* model, since its main goal is not describing or explaining patterns in how actual agents are disposed to make decisions.[6]

3 Putting the 'normative' in normative modelling

The distinction between normative and descriptive modelling is intuitive. But what exactly does it consist in? As I see it, the crucial difference between normative and descriptive models is that normative models are centrally connected with what one ought to do. Descriptive models are not. Some advocates of normative modelling have discussed this connection. For instance, Colyvan (2013, 1347) expresses a version of this view in terms of normative models delivering good advice. He writes:

> We need to pay attention to the intended purposes of the model and the level of detail it is supposed to provide... The primary purpose of normative models seems to be to deliver good advice about decisions, inferences, the structure of beliefs and the like—at some appropriate level of abstraction.

The view that normative models are centrally connected to some sort of 'ought' captures an idea in fairly coarse-grained detail. In this section, I make it a little more precise and defend it from a possible objection. Let's consider three more specific ways of specifying this connection: a *weak connection*, a *strong connection*, and a *moderate connection*. I argue that we should accept the *moderate connection*. It is informative enough to capture, at a certain level of generality, the connection between normative models and some sort of 'ought',

[6] Although I have emphasised the distinction between normative and descriptive models, one might wonder whether the difference is quite so sharp. Perhaps models lie on a continuum. On this view, 'normative' and 'descriptive' are best seen as labels for models that cluster toward either end of the spectrum, with many shades of grey in between. This is an interesting suggestion that I do not wish to dismiss. One can reasonably interpret Titelbaum's (2025) and Roussos's (2025) views as consistent with this picture: the more frequently a model is put to a normative use (Roussos), or the more often it targets normative facts (Titelbaum), the further toward the 'normative' end of the spectrum it sits. For my purposes in this paper, the relevant point is that we can draw a meaningful distinction between normative and descriptive models, regardless of whether one thinks of this as a binary distinction or as describing clusters on an underlying spectrum. For simplicity, I have presented this as a clear-cut distinction, but I am open to the possibility that a more nuanced picture underlies the differences I emphasise. In Section 3, I develop the normative/descriptive distinction in a way that is also compatible with a spectrum-based view. I am grateful to an anonymous referee for helpful comments on this point.

while avoiding implausible implications that arise if we think about the connection in either too *strong* a way or too *weak* a way.

On a *weak connection*, the relationship between normative models and 'ought' is roughly the following: normative models can deliver normative guidance about what you ought to do. This seems much too weak to me. Descriptive models can also deliver advice about what you ought to do, in some sense. Depending on some other conditions being in place, descriptive models can deliver advice about all sorts of things, from how governments ought to regulate carbon emissions to how individuals ought to save for the future. All that it takes for this connection to obtain is that a model sometimes has practical implications, which scientific or economic models often do. If normative modelling can be distinguished from descriptive modelling in virtue of its connection with an 'ought,' then the relationship must be tighter than this.

On a *strong connection*, the relationship between normative models and 'ought' is roughly the following: normative models always deliver normative guidance about what you ought to do. This seems much too strong to me. Delivering advice is not the only purpose of a normative model, there are plenty of other ways to use them. For instance, they can uncover results that are interesting or valuable in their own right, or valuable for future inquiry but would rarely connect with what one ought to do in any significant way. There are also other *normative* purposes apart from delivering advice. Consider standard decision theory again. This is a paradigmatic example of a normative model, but it is not the case that the only normative role it can play is to deliver advice about what one rationally ought to do.

For instance, there is the *interpretative* way of understanding standard decision theory, most prominently associated with Ramsey (1926).[7] This approach treats decision theory as a way to make sense of an agent's beliefs and desires. It takes an agent's preferences as a starting point and asks what beliefs and desires they must have if we are to interpret them as rational. Note that this is not a *descriptive* use of decision theory (Buchak 2013, 85). The goal of the interpretive use of decision theory is not to describe the cognitive patterns of actual agents or predict how they will choose. Rather, it is to present a coherent picture of an agent's mental states on the assumption, as standard decision theory holds, that preferences are rationally linked with beliefs and desires. It is still doing normative decision theory; it is just putting the standard model to

[7]See also Buchak (2013, 83-85), Lewis (1974) and Davidson (1973).

a different kind of normative use than providing advice about what one ought to do. The upshot of the interpretive use of decision theory is that one can uncover something that is interesting in its own right; one can describe an agent's beliefs and desires, and one can provide a *rational* explanation of them given the agent's preferences. Since none of this constitutes normative guidance, and since the *strong connection* holds that normative models always deliver normative guidance, the *strong connection* rules out putting decision theory to this use. There is no need to insist on such a narrow range of purposes, and I take it that the case of standard decision theory shows that a normative model can have multiple normative uses, such as guidance *and* interpretation.[8]

I favour a *moderate connection* between normative models and 'ought' along the following lines: normative models naturally and usually deliver normative guidance about what you ought to do. Why accept the *moderate connection*? I think it captures the spirit of the intuitive idea that is often advocated by normative modellers, such as Colyvan (2013, 1347), and it avoids the extremes of the other two views.

The degree of modal robustness present in the *moderate connection* is also supported by two features of normative modelling highlighted by Titelbaum (2025) and Roussos (2025). Recall that on Titelbaum's view, normative models are a different *kind* of model to scientific models; in particular, they are models that target normative data, such as evaluations of actions and attitudes. On Roussos' view, normative models are not a different kind of model, rather they are a model put towards a normative purpose, such as action-guidance. On *both* Titelbaum's and Roussos' way of distinguishing normative models from descriptive ones, the connection with some sense of 'ought' is central—Titelbaum's emphasis on normative data includes questions about what one ought to do and believe, and Roussos' emphasis on normative purposes includes action-guidance. But neither writer takes it that advice *exhausts* the normative data a model might target or the normative purposes a model might be used for. Both include evaluation, which can sometimes come apart from action-guidance (e.g., Smart 1973), and both include investigating normative concepts, which again might sometimes come apart from questions about what one ought to believe or do (e.g., Sepielli 2012). Both accounts draw the distinction between normative and descriptive modelling in a way that supports the *moderate connection*. This way of refining the rough idea that normative mod-

[8]I am grateful to an anonymous referee for helpful comments that improved this section.

els are connected with some kind of 'ought' is both intuitively attractive and is supported by several considerations from the nature and aims of normative modelling.[9]

So far, I have emphasised that the *moderate connection* precisifies the connection between normative modelling and some sense of 'ought' that advocates of normative modelling often gesture towards. But I think that it also makes the very idea of a connection more defensible. Consider the following objection:

> *Objection: Although normative models can provide guidance, this is not the best way to distinguish them from descriptive models. They have a variety of practical uses, and insisting on a central connection between normative models and some kind of 'ought' distorts the benefits that they might provide.*

This is the kind of objection that someone sympathetic to a view like Roussos' (2025) might want to press. On his view, providing advice about what actual agents ought to believe, or how they ought to behave, is too narrow a way of conceptualising what distinguishes a normative model from a descriptive one. He takes it that a normative model is any model that is put to a normative purpose, and *this* includes more than just guidance (Roussos 2025, 431).

Although Roussos does not mention this example himself, I take it that much formal work on the social evolution of norms supports this way of thinking about the goals of normative modelling. Philosophical literature on this topic develops models aimed at understanding how different kinds of norms have evolved, often modelling the introduction of these norms as a rational response to a coordination problem or as an answer to a practical need (e.g., Lewis 1969, Skyrms 1996).

It does not seem apt to classify the models developed in this kind of work as descriptive models, since they hope to capture how a rational group of agents *ought* to collectively behave in the face of a coordination problem, rather than

[9] *The moderate connection* heads off another challenge. One might worry that I slide too quickly from speaking about an 'ought' to speaking about 'advice' or 'normative guidance.' There are different kinds of 'oughts.' Not all of them are especially guidance-giving; and even those that are, can deliver advice that is more or less feasible (e.g., Southwood 2016, Jackson and Smith 2016). Isn't there an important difference between the claim that normative models are connected with some kind of 'ought', and the claim that normative models deliver normative guidance? The way of spelling out the connection between normative modelling and 'ought' that I favour can respond to this challenge. According to *the moderate connection*, normative models naturally and usually deliver normative guidance. This interpretation allows one to focus on the connection between normative models and normative guidance, while avoiding the misunderstanding that 'oughts' always issue feasible advice.

how a group in fact behaves. But, at the same time, it also does not seem accurate to think that these models have the main goal of delivering normative guidance about how we ought to reason or behave. The chief benefit of these models lies in *vindicating* some social practice as rational or beneficial, or *explaining* the functional role of features of our social life. There is something odd about insisting that the main benefit of social evolutionary models of normative practices is to provide us with advice about how we ought to cooperate and bargain when faced with a coordination problem.

The *moderate connection* suggests two lines of response to this objection. First, it can accommodate much of this challenge in a way that leaves its core idea intact. Although I think that normative models are centrally connected with normative guidance, I have rejected thinking about this in terms of a *strong connection* whereby 'normative models always deliver normative guidance about what you ought to do.' This allows me to accommodate several points that motivate this objection. I can accept Roussos' general point that delivering advice is not the only purpose of a normative model. And I can accept the lesson drawn from the example about the social evolution of norms: sometimes normative models are perfectly good models yet fail to deliver useful advice. Thinking of the relationship between normative models and advice in terms of a *strong connection* would lead one to reject both points, but a *moderate connection* can accommodate normative models having a variety of purposes.

But perhaps someone pressing this objection might want to see normative guidance assigned an even smaller role. Given the plurality of normative purposes that I emphasised in my previous reply, they might point out that, by my own lights, downgrading the importance of normative guidance is compatible with normative models being properly normative.

My second line of reply addresses this more specific point. I think this objection leaves us with an unfamiliar picture of normativity. The *moderate connection* is supported by quite general considerations about normativity *simpliciter*. Most normative mechanisms are centrally connected with some kind of 'ought.' It is standard to think that this is one of the fundamental ways in which a mechanism might be normative in the first place. Philosophers who accept that part of a theoretical system is normative—for example, logic, prudence, or morality—tend to accept that its normativity partly consists in its connection with an 'ought.' To think that a philosophical account of logic or justice is normative is usually, *inter alia*, to think that the account is connected with how we ought to reason or arrange a society. The *moderate connection*

piggybacks on this familiar picture of normativity, emphasising a similar connection that one tends to find across the normative landscape. Were one to weaken the connection with normative guidance further, it would introduce a significant asymmetry between normative modelling and other branches of normative philosophy; a picture on which a mechanism is properly thought of as normative but is not centrally connected with some kind of 'ought.' And I think this view would need to be more thoroughly motivated.

At this point, it also matters how *severe* the downgrading is that is being proposed by the objector. For those who wish to downgrade normative guidance to the point where it plays a much smaller role than in other normative domains, the above reply seems apt. But some may not wish to go quite this far, while still insisting that I am overvaluing the importance of normative guidance. An example of this kind of view can be found in writers such as Christensen (2004, 143-178) who claim that the primary normative purpose of formal epistemology is to describe the ideal of rational belief. Describing ideal performance is a legitimate normative purpose, even though, on this view, looking to a formal model for cognitive advice is unlikely to be particularly helpful (Christensen 2004, 167-168). However, it is worth noticing that many writers who advocate views like this take it that the normative models of formal epistemology describe the ideal that we *ought* to strive towards. The closer we get, the better we have done (e.g., Christensen 2004, 165-167; Smithies 2015, 2781). The worry, it seems to me, that views like this raise is not that normative models aren't importantly connected with some kind of 'ought.' Rather, the challenge is whether the model can generate useful guidance—advice that might help an actual agent to approximate the ideal—given that the ideal they describe is often far beyond human attainability. In Section 4 below, I take up this topic, arguing that normative models can give advice.[10]

4 Normative models can give advice

So far, I have put forward one main claim: one role of normative models is to deliver normative guidance. This involved describing the enterprise of normative modelling and arguing that part of its difference from descriptive modelling is that it is centrally connected with what one ought to do or believe (Section 2). I have also argued that this connection can be made more precise

[10] I am grateful to an anonymous referee for helpful comments on this point.

and defensible if we characterise it in terms of a *moderate connection*: normative models naturally and usually deliver guidance about what one ought to do or believe (Section 3).

Of course, *should* and *can* are two different matters, and one might wonder: even if normative models should deliver guidance, *can* they? 'Oughts' are notoriously complicated in many areas of philosophy. An 'ought' derived from a model just adds more complexity. For this paper's proposal to be plausible, we need a clearer account of how a normative model might deliver normative guidance.

In this section, I show that normative models can generate normative guidance. Drawing on Titelbaum (2012), I describe a promising framework for thinking about the connection between normative models and normative guidance. One attractive feature of Titelbaum's framework is that it is amenable to expansion, and so I also discuss how we can enrich this account. Reflecting on Titelbaum's framework also helps to clarify the nature of the advice that a normative model might provide. I conclude this section by addressing two objections that one might have to this general picture.

4.1 Titelbaum on models and norms

Titelbaum (2012, 11-27) provides a framework for thinking about the relationship between formal models and norms. It will help to approach this discussion with a concrete example in mind of a normative model and an epistemic norm.[11] Let's take a standard Bayesian conditionalization model like the below:

> **CONDITIONALIZATION MODEL:**
> $Cr_j(H) = Cr_i(H|E)$

CONDITIONALIZATION MODEL is a model of how credences change after a learning event. Cr_i is an agent's credence distribution before they learn a proposition E, and Cr_j is the credence distribution they ought to have after learning E. The model motivates a variety of norms of rationality; paradigmatically, the requirement that rational agents update their attitudes in this

[11] It bears emphasising that the relationship between models and norms is not just a problem for Bayesianism or formal epistemology. For instance, the problem could just have easily been specified in terms of reasoning norms and formal logic. See Titelbaum (2012, 12-14) for a presentation of the problem in these terms.

way. When a rational agent learns a proposition, their new credence distribution should be their prior credence distribution conditional on what they just learned. The norm is usually represented like this:

> **RULE OF CONDITIONALIZATION:**
> Rationality requires that, for all propositions H, $Cr_j(H) = Cr_i(H|E)$

And of course, the model has several variations and is used to motivate a range of other epistemic norms.[12] At first pass, it might seem that there is barely any difference between conditionalization as a model of a learning event and an epistemic norm like the **RULE OF CONDITIONALIZATION**. If that is the case, norms are transparent to models; we can read an epistemic requirement straight off a normative model. What is so bad about that?

Two things. First, this is not an adequate way to approach modelling in general. All models need to be interpreted. To do this, we need a 'key' that specifies how elements of the model are to be mapped onto the target system (Frigg and Nguyen 2016, 2023). Second, the *moderate connection* holds that normative models naturally and usually deliver normative guidance about what one ought to do. Part of the case I wish to explore for this proposal involves clarifying how normative models can play this role. For many pairs of models and norms, the guidance that a model might provide cannot be gleaned in as straightforward a way as in the case of conditionalization. Indeed, Titelbaum (2025, 55-57) and others have stressed that even in the case of conditionalization, the relationship between the normative model and the advice is not as straightforward as it might seem. One benefit of Titelbaum's account of a key for mapping models onto norms is that it clarifies how a normative model can provide guidance.

Let's turn to Titelbaum's framework now. He takes it that modelling is comprised of four main elements: *a modelling framework*, *the model itself*, *bridge principles*, and *modelling rules*. We do not need to spend much time on the first two. The model itself is the formal object—usually a piece of mathematics like a set of equations—used to capture some features of the target. The modelling framework is the background apparatus one needs to construct multiple models of a target. It includes things like what features of the target are most salient, systematic constraints common to every model we might build of the target, and the mathematics that is best suited to capturing the target's features.

[12] For a helpful overview, see Meacham (2015) and Pettigrew (2020).

I will set aside what Titelbaum has to say about these in order to focus on the other two elements—*bridge principles* and *modelling rules*—both of which help us understand how to read advice off a normative model.

Bridge principles relate a formal model to its target.[13] The way they do so is quite subtle, so it is worth quickly setting out the account. On Titelbaum's view, there is a third relatum that connects a formal model with a system of norms, namely a philosophical concept. Bridge principles connect the three relata. We can represent the relationship between a model, a philosophical concept, and a set of norms schematically as in Figure 1 below.

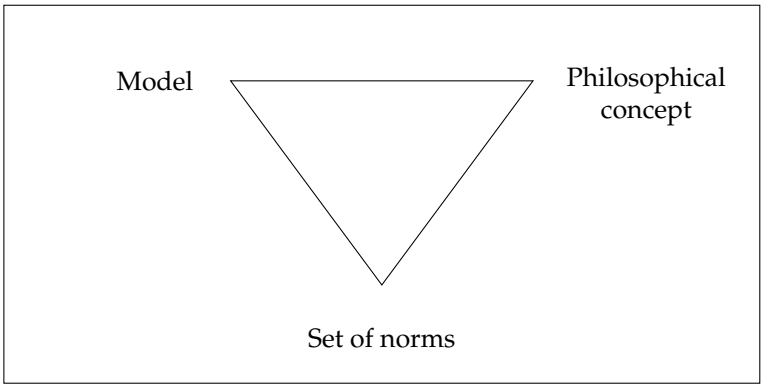

Figure 1: *Titelbaum's Triangle*

In Titelbaum's Triangle (Figure 1), bridge principles are represented by the lines that connect the three relata; bridge principles are paths that allow you to move from a normative model to a philosophical concept to a set of norms. A couple of examples will help make this clearer. Take deductive logic. Here, the *formal model* is a deductive system like classical logic, the *philosophical concept* is logical consequence, and the *norms* are norms of reasoning. Filling in the above scheme, Titelbaum's Triangle for norms of reasoning looks like Figure 2 below.

Or take an example from Bayesian epistemology. In this case, the *formal model* is a probability function, the *philosophical concept* is subjective probability, and the *set of norms* are rational requirements on credences. Filling in Figure 1, Titelbaum's Triangle for Bayesian epistemology looks like Figure 3 below.

At first pass, this way of setting out the relationship between models and norms might seem unnecessarily complex; it adds in the notion of a philosoph-

[13]Titelbaum notes that his use of the term 'bridge principle' is quite general (Titelbaum 2012, 17 n. 3). It is worth emphasising that he takes this term from John MacFarlane's work on the normativity of logic (MacFarlane 2004), rather than from the philosophy of science.

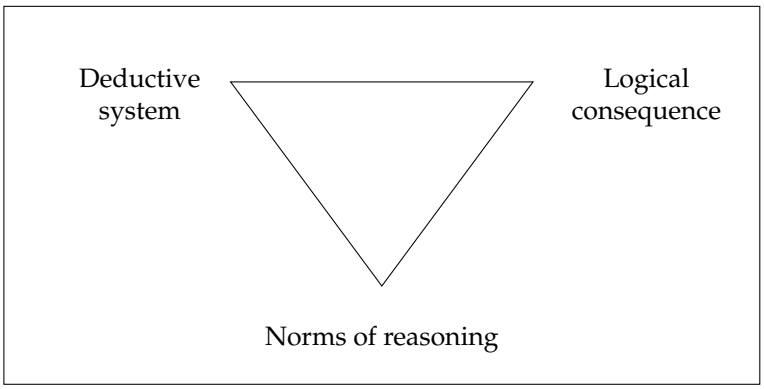

Figure 2: *Titelbaum's Triangle for Reasoning*

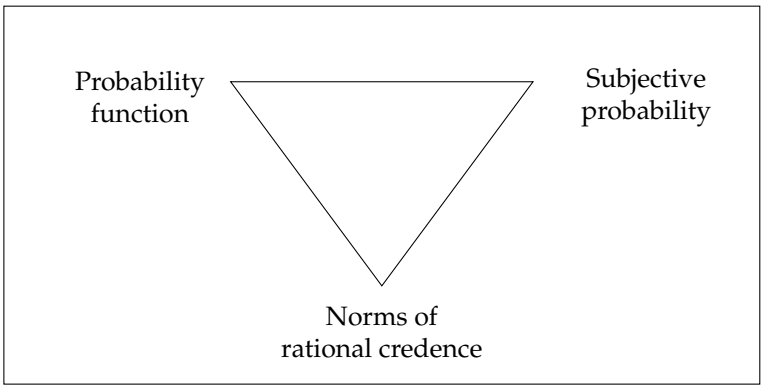

Figure 3: *Titelbaum's Triangle for Bayesian Epistemology*

ical concept and involves several bridge principles that connect the relata. This multiplies layers and mechanisms between the normative model and set of norms. But the benefit of Titelbaum's Triangle is that it spotlights clear choice points when one considers how a model might relate to a set of norms. One advantage of identifying these elements is that it demonstrates that just because a feature appears in a model does *not* entail that it must appear in the corresponding set of norms. More precisely, it allows one to identify the principles that take a feature from one relatum and move it across to another. For instance, those who accept classical logic standardly accept the following two principles:

1. If $x \vdash y$ in a classical deductive system, then y follows from x.

2. If y follows from x, then it is permissible to reason from x to y.

Titelbaum suggests that these are the bridge principles in Figure 2 above. Principle 1 connects the formal model of classical logic with the philosophical concept of logical consequence, since the antecedent is a feature of a formal system and the consequent is a philosophical concept—in this case, the formal feature is the turnstile, and the philosophical concept is logical consequence. In the same way, Principle 2 connects the philosophical concept of logical consequence with the norm of a permissible inference. Titelbaum's way of thinking about the structure of norms and models allows us to do two things. First, we can *identify* when a feature is part of a formal model but is not part of a set of norms. And second, we can *vindicate* this result by specifying the bridge principles that explain why a feature of a model does not carry over to a set of norms.

Let's look at the second element of Titelbaum's framework. A *modelling rule* tells you when to take a model's advice. One such modelling rule is that one ought to consider the domain of applicability (Titelbaum 2012, 22). Modelling rules provide a guide for determining when to use a particular model and when to take its advice. It does not contribute to what the model's advice is, but it tells you when those verdicts are to be trusted (Titelbaum 2012, 22). An important upshot here is that just because a model's advice is not to be trusted outside its domain of applicability should not diminish our confidence in the advice it provides *within* that domain (Titelbaum 2012, 23). This is a point that Roussos (2025) also emphasises.

With this account of *bridge principles* and *modelling rules* set out, let's return to our original example of conditionalization. Titelbaum's framework has a number of upshots for how we think about the connection between conditionalization as a formal model and a norm of rationality. It is worth highlighting a number of these in order to illustrate how the framework can be put to use. Here are four lessons we can derive from it.

First, the framework highlights that the notion of a 'rational requirement' is a philosophical concept that serves as an intermediary relatum between the formal model and the updating norm—notice that it does not appear in **CONDITIONALIZATION MODEL**, but it does appear in the **RULE OF CONDITIONALIZATION**. The notion of a rational requirement is the subject of philosophical work in its own right (e.g., Broome 2013, Kiesewetter 2017). Titelbaum's framework allows that some concerns about conditionalization as a norm might be

about the philosophical concept of a rational requirement.

Second, as Titelbaum emphasizes elsewhere (2025, 53-54), his framework shows that attitude changes that are *not* learning events—that is, forming a credence less than 1 towards E—lie outside the domain of applicability of the model. A common objection to conditionalization as an epistemic norm is that it seems to require one to be certain of new evidence and to never reduce one's confidence in a certainty, both of which present an unfamiliar picture of how we respond to experience. Titelbaum's framework provides a response to these challenges, since many of the cases lie outside of the domain of applicability of the attitudinal changes to which the model can be appropriately applied.

Third, the framework highlights the need for bridge principles. For instance, $Cr_j(H)$ appears in both the model and the norm. In the model, its interpretation is straightforward; it represents a property of a probability space. But norms are not purely formal objects, so its interpretation is less clear. Without plausible bridge principles, conditionalization is susceptible to standard objections that it is too far removed from the mental lives of ordinary thinkers. Bridge principles can address these worries, say, by establishing that in the RULE OF CONDITIONALIZATION $Cr_j(H)$ relates to a feature of an agent's doxastic attitudes, such as their total credal state, or to their confidence in H, or to their disposition to act on the basis of H, and so forth.

Fourth, the framework insulates the mathematics from the epistemic requirements. Depending on one's bridge principles, and what the philosophical notion of 'rational requirement' amounts to, the content of the formal model and set of norms will be different. I take it this is a feature, not a bug; it is the framework working as it is supposed to, ensuring that not 'everything appearing in our formal model's equations must appear in the agent's head,' as Titelbaum (2012, 17) puts it. This heads off the worry that choices about what norms one endorses will lead to *ad hoc* revisions to firmly established formal systems, such as the probability calculus.

Of course, many of these points are controversial. I mention them here not to advocate for them as substantial views, but just to illustrate how fruitful Titelbaum's framework is when applied to a particular issue.

Titelbaum's framework is also amenable to expansion. Features of normative models emphasized by other writers can be incorporated into the framework. For instance, two lessons from modelling in economics can be neatly accommodated as additional *modelling rules*.

First, a number of writers emphasise that economic models are only useful

if their 'critical assumptions' are true (Gibbard and Varian 1978, 671; Rodrik 2015, 27-29, 95-98). Critical assumptions are those whose closeness to an actual scenario substantially affects the results of the model. For example, the assumption of a perfectly competitive marketplace will often make no difference to the advice an economic model generates. But the assumption of a monopolized marketplace will make a difference depending on whether this assumption is borne out in the targets. Many assumptions are not critical in this respect, but for the model to be useful for a particular purpose, the critical assumptions have to track reality to a sufficient degree. We can think of this point as identifying an additional modelling rule; that is, a parameter that determines when the advice delivered by a normative model is suitable.

Second, economists (and others) standardly rely on a library of models, each of which provides advice in different scenarios (Rodrik 2015, 20). They treat models as providing conditional advice; roughly, if some circumstances C obtain, then model M provides good advice, but if different circumstances obtain, then M does not provide good advice and one ought to consult a different model from the library. This is a natural consequence of modelling methodology: since models do not aim at being the only true model, and since they partially illuminate complex targets, it is to be expected that models multiply without replacing one another (Rodrik 2015, 67). Again, we can think of the range of available models as suggesting a further parameter that determines how one ought to apply a model. In addition to considering the domain of applicability, and whether the model's critical assumptions hold, we also ought to consider whether one has selected the best available model for one's purposes.

To conclude, it is worth emphasising that Titelbaum's framework identifies an additional reason to be satisfied with the *moderate connection*. In his framework, the advice comes *from* the normative model; Titelbaum's Triangle begins with the formal model and ends with a set of norms. This provides some support for the 'naturalness' feature of the *moderate connection*. Titelbaum's Triangle provides grounds to think that there is a natural connection between normative models and normative guidance, even if the guidance is filtered through several other mechanisms, such as a *philosophical concept* and some *bridge principles*. But, nonetheless, the framework shows that the advice that ends up in the set of norms was contained in the formal model in the first place.

4.2 Objections and replies

In this section, I have argued that normative models can provide normative guidance, and I have drawn on Titelbaum (2012) to describe how they might do so. The resulting picture is complicated, so it is worth considering a couple of possible objections that one might have. The first objection addresses this *complexity*, and the second addresses the *nature* of the advice that the framework delivers.

> *Objection: Although delivering advice might be one feature of a good normative model, another is simplicity. The priority I give to advice is not compatible with simplicity—the more we expect a normative model to provide advice, the messier it will become.*

All models harness the power of simplicity; it is one of their chief virtues. To make the case that normative models can deliver guidance, I have drawn on a complex framework. Simplicity is not the *only* theoretical virtue of a model—models should be no simpler than they need to be. Nevertheless, this objection presses the point that my account gets the balance wrong: it downgrades a central desideratum of modelling and upgrades a peripheral one. This challenge can be met in three ways.

My first reply appeals to a standard feature of modelling; namely that we should not expect any model to be the one true model of a phenomenon. To adequately capture a target system—including systems in normative domains—we need a library of models.[14] Different models in the library will have different properties which make them apt for different tasks. If we find that some normative model is too messy for some purpose, one response would be to see if the model can be improved by making it simpler. But this is not the only approach the methodology affords. Another option that fits well with modelling in general is to develop different models for different purposes. Having a range of models relieves us from needing to insist that all models in our library be as simple as one another, or meet some threshold of simplicity. This is a feature that differentiates modelling from theory building, and an advantage of modelling is that it can shuttle between a multiplicity of models tailored to a variety of purposes. The modeller has more options here than the theoretician—they might explore simplifying their model (or improving

[14] Daniel Greco's (2023) recent defence of 'modest modelling' in epistemology also emphasises the importance of looking to a library of models for a variety of purposes.

it along some other axis) in a way that preserves its current theoretical utility, or they might look elsewhere for a simpler model when more simplicity is needed.

Second, depending on how broad a view one takes of the notion of simplicity as a virtue, I think my view promotes simplicity in our normative epistemic practices. For instance, one application of my view is that, if one accepts Titelbaum's suggestion that the domain of applicability for the **RULE OF CONDITIONALIZATION** is all *learning* events rather than all updating events, then conditionalization is not a norm that provides advice right across the spectrum of various kinds of attitudinal change.[15] This result leads to greater simplicity in some of our everyday epistemic practices, since the set of attitudes governed by this norm is reduced in size.

Third, a feature of Titelbaum's Triangle provides an additional line of response to this objection. The framework allows us to identify the level at which this complexity emerges. It does *not* emerge at the formal model vertex. In Titelbaum's Triangle, complexity arises at the other levels; that is, from taking a formal model and connecting it with a philosophical concept, and a set of bridge principles, and applying it to a concrete case based on the parameters picked out by a set of modelling rules, and so on. The model itself is insulated from the complexity generated by other parts of the framework.

On my view, the models themselves do not become any less simple, and it is not obvious to me that the broader picture of how models are put to normative uses sacrifices simplicity to a concerning degree.

> *Objection: Demonstrating that normative models can generate mere advice is inadequate for my purposes. Advice can be apt but useless. For my argument to go through, I need to show that normative models can provide advice that is useful or followable.*

Some features of Titelbaum's Triangle also help to address this challenge. On this way of thinking about the relationship between models and norms, there are mechanisms that help to promote positive features in the set of norms that are derived from the model. I think these include some features that norms which provide good quality advice would have. The mechanisms I have in mind are the bridge principles and the modelling rules. For instance, two of the modelling rules I discussed above capture considerations about a model's

[15] Of course, some philosophers have proposed ways of adapting conditionalization so that it can be put to a broader use along these lines, most prominently Jeffrey (1983).

domain of applicability and its critical assumptions. These modelling rules create parameters that, when taken together with other modelling rules, serve to pick out the particular cases from a range of possible cases for which the norms are relevant. These rules ensure that a model provides *conditional* advice, advice that can be trusted conditional on certain background conditions holding, such as an appropriate target of the norm falling within the domain of applicability, or the model's critical assumptions being borne out in a particular concrete case. This goes some way towards ensuring that the norms provide appropriate advice.

I do not want to over emphasise the advantages of Titelbaum's framework; it is not an infallible method for reading good advice off a normative model. But it does include within its apparatus features that support this. By requiring one to specify the bridge principles that connect a model with a set of norms, and the modelling rules that determine under what circumstance the norms apply to a case, the framework goes some way towards ensuring that the advice that is generated will have some of the properties of good advice, such as not being too abstract, trivial, or useless. And again, I think this framework provides the model builder with advantages compared to the theoretician faced with the same challenge. Titelbaum's Triangle makes explicit important choice points in determining how a normative model generates advice in a way that often is not explicit when one considers how a normative theory does the same.

We may also partly respond to this challenge with a partners-in-guilt defence. This question essentially raises the puzzle of how an agent is guided by a norm. This a significant topic in its own right, one that has been wrestled with across different domains and eras of philosophy.[16] Whether a piece of advice is generated by a normative model, a philosophical theory, or by folk wisdom and ordinary morality, vexed questions arise about the conditions that need to obtain for it to be the case that the norm provided an agent with good guidance. I do not want to dismiss these philosophical problems. But I do not think they present a special challenge to my view.

[16]For a recent and detailed discussion of the 'usability' problem in ethics, see Smith (2018) and Jackson and Smith (2016).

5 Conclusion

In this paper, I argued that a particular way of defending highly demanding epistemic norms—what I have called *the modelling defence*—fails. But its failure is not really bad news. It reveals reasons internal to the enterprise of normative modelling that can be used to build an attractive picture of some of the goals of formal epistemology and, more broadly, model building in philosophy. On this picture, there is a natural connection between normative models and normative guidance such that it is appropriate to expect formal epistemological models to deliver some useful advice. This is quite a different picture than one often encounters, especially among philosophers who might have their doubts about the usefulness of formal work in epistemology. But it deserves consideration, especially in light of a larger interest in finding ways past received distinctions between traditional/formal and ideal/non-ideal epistemology.

My argument for this view has been in two parts. First, I argued that drawing a distinction between normative and descriptive modelling involves connecting normative modelling with some kind of an 'ought' and I argued for a *moderate* interpretation of this connection (Section 3). On this view, one key difference between normative and descriptive models is that normative models naturally and usually deliver normative guidance about what one ought to do, whereas descriptive ones do not. Second, drawing on a framework due to Titelbaum (2012), I argued that normative models can deliver on this promise (Section 4). Taken together, the lines of argument make the case that normative models *should* and *can* deliver normative guidance. In other words, we can have formal epistemology without demandingness.

To conclude, let's return to the issue with which I began to see how the view I have developed can be put to work. Take a common demandingness objection raised against a norm of formal epistemology; for instance, the objection that conditionalization is extremely computationally demanding and it is unrealistic to expect actual agents to satisfy it. In response, my view can deliver good news for both the objector—the person who worries that conditionalization is too demanding—and the formal epistemologist committed to its importance.

To the objector, my view replies that norms such as this can and should deliver normative guidance. Moreover, it encourages them to temper their expectations by appreciating that the type of advice that it is reasonable to expect from a normative model is often partial, conditional, and the result of careful

application via a framework such as Titelbaum's. The upshot is that if a norm like conditionalization rarely provides useful advice—taking into consideration the application framework described in Section 4 and the incomplete nature of the advice—then, on my view, it runs afoul of the goals of normative modelling, and the demandingness objection stands. But my view also has good news for the formal epistemologist by setting out several promising lines of reply. They might, for example, show that the model yields advice that is more domain-specific than the objection presumes, or that another model in the library yields better advice. On my view, normative models naturally and usually deliver normative guidance, and they do so through several mechanisms that connect a model with some advice. When a norm faces a demandingness objection, my view provides the formal epistemologist with options for satisfying the objection without revising the model itself.[17]

[17] I am grateful to an anonymous referee for helpful comments on this point.

References

Appiah, Kwame Anthony. 2017. *As If: Idealizations and Ideals*. Harvard University Press.

Broome, John. 1991. *Weighing Goods*. Wiley-Blackwell.

———. 2013. *Rationality Through Reasoning*. Wiley-Blackwell.

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

Christensen, David. 2004. *Putting Logic in Its Place: Formal Constraints on Rational Belief*. Oxford University Press.

Colyvan, Mark. 2008. 'Is Probability the Only Coherent Approach to Uncertainty?' *Risk Analysis* vol. 28, no. 3: 645-652.

———. 2013. 'Idealisations in Normative Models.' *Synthese* vol. 190, no. 8: 1337-1350.

———. 2014. 'The Undeniable Effectiveness of Mathematics in the Special Sciences.' In *New Directions in the Philosophy of Science*, eds., M. C. Galavotti, et al. Springer, 63-73.

Davidson, Donald. 1973. 'Radical Interpretation.' *Dialectica* vol. 27, no. 3-4: 313-328.

Earman, John and John Roberts. 1999. '*Ceteris Paribus*, There is no Problem of Provisos.' *Synthese* vol. 118, no. 3: 439-478.

Earman, John, John Roberts, and Sheldon Smith. 2002. '*Ceteris Paribus* Lost.' *Erkenntnis* vol. 52, no. 3: 281-301.

Friedman, Jane. 2018. 'Junk Beliefs and Interest-Driven Epistemology.' *Philosophy and Phenomenological Research* vol. 97, no. 3: 568-583.

Frigg, Roman and James Nguyen. 2016. 'The Fiction View of Models Reloaded.' *The Monist* vol. 99, no. 3: 225-242.

———. 2023. 'Maps, Models, and Representation.' In *Scientific Understanding and Representation: Modeling in the Physical Sciences*, eds., Insa Lawer, Kareem Khalifa, and Elay Shech. Routledge, 261-279.

Gettier, Edmund. 1963. 'Is Justified True Belief Knowledge?' *Analysis* vol. 23, no. 6: 121-123.

Gibbard, Allan and Hal R. Varian. 1978. 'Economic Models.' *The Journal of Philosophy* vol. 75, no. 11: 664-677.

Godfrey-Smith, Peter. 2007. 'The Strategy of Model-based Science.' *Biology & Philosophy* vol. 21, no. 5: 725-740.

———. 2009. 'Models and Fictions in Science.' *Philosophical Studies* vol. 143, no. 1: 101-116.

Greco, Daniel. 2023. *Idealization in Epistemology: A Modest Modeling Approach.* Oxford University Press.

Hintikka, Jaakko. 1962. *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Cornell University Press.

Ismael, Jenann. 2016. 'A Philosopher of Science Looks at Idealizations in Political Theory.' *Social Philosophy and Policy* vol. 33, no. 1-2: 11-31.

Jackson, Frank and Michael Smith. 2016. 'The Implementation Problem for Deontology.' In *Weighing Reasons*, eds., Errol Lord and Barry Maguire. Oxford University Press, 279-293.

Jeffrey, Richard C. 1983. *The Logic of Decision*. Second edition. University of Chicago Press.

Kahneman, Daniel and Amos Tversky. 1979. 'Prospect Theory: An Analysis of Decision under Risk.' *Econometrica* vol. 47, no. 2: 263-291.

Keefe, Rosanna. 2000. *Theories of Vagueness*. Cambridge University Press.

Kiesewetter, Benjamin. 2017. *The Normativity of Rationality*. Oxford University Press.

Lewis, David. 1969. *Convention*. Harvard University Press.

———. 1974. 'Radical Interpretation.' *Synthese* vol. 23, no. 3-4: 331-344.

MacFarlane, John. 2004. 'In What Sense (If Any) is Logic Normative for Thought?' Unpublished. https://johnmacfarlane.net/normativity_of_logic.pdf

Meacham, Christopher J. G. 2015. 'Understanding Conditionalization.' *Canadian Journal of Philosophy* vol. 45, no. 5: 767-797.

Parker, Wendy S. 2009. 'Confirmation and Adequacy-for-Purpose in Climate Modelling.' *Proceedings of the Aristotelian Society, Supplementary Volumes* vol. 83, no. 1: 233-249.

Pettigrew, Richard. 2020. 'What is Conditionalization, and Why Should We Do It?' *Philosophical Studies* vol. 177, no. 11: 3427-3463.

Ramsey, Frank. 1926. 'Truth and Probability.' In *The Foundations of Mathematics and Other Logical Essays*, ed., R. B. Braithwaite. Routledge and Kegan Paul, 156-198.

Rodrik, Dani. 2015. *Economics Rules: The Rights and Wrongs of the Dismal Science*. Norton.

Roussos, Joe. 2022. 'Modelling in Normative Ethics.' *Ethical Theory and Moral Practice* vol. 25, no. 5: 865-889.

———. 2025. 'Normative Formal Epistemology as Modelling.' *The British Journal for the Philosophy of Science* vol. 76, no. 2: 421-448.

Sepielli, Andrew. 2012. 'Subjective Normativity and Action Guidance.' In *Oxford Studies in Normative Ethics* vol. 2, ed., Mark Timmons. Oxford University Press, 45-73.

Skyrms, Brian. 1996. *Evolution of the Social Contract*. Cambridge University Press.

Smart, J. J. C. 1973. 'An Outline of a System of Utilitarian Ethics.' In *Utilitarianism: For and Against*, eds., J. J. C. Smart and Bernard Williams. Cambridge University Press, 3-74.

Smith, Holly M. 2018. *Making Morality Work*. Oxford University Press.

Smithies, Declan. 2015. 'Ideal Rationality and Logical Omniscience.' *Synthese* vol. 192, no. 9: 2769-2793.

Southwood, Nicholas. 2016. 'Does "Ought" Imply "Feasible"?' *Philosophy & Public Affairs* vol. 44, no. 1: 7-45.

Stalnaker, Robert. 1978/2002. 'Assertion.' In *Formal Semantics: The Essential Readings*, eds., Paul H. Portner and Barbara. H. Partee. Blackwell, 147-161.

Titelbaum, Michael G. 2025. 'Normative Modeling.' In *Methods in Analytic Philosophy: A Primer and Guide*, eds., J. Horvath, S. Koch, and M. Titelbaum. The PhilPapers Foundation, 41-66.

———. 2012. *Quitting Certainties: A Bayesian Framework Modeling Degrees of Belief*. Oxford University Press.

Tversky, Amos and Daniel Kahneman. 1992. 'Advances in Prospect Theory: Cumulative Representation of Uncertainty.' *Journal of Risk and Uncertainty* vol. 5, no. 4: 297-323.

Weisberg, Michael. 2007. 'Three Kinds of Idealization.' *The Journal of Philosophy* vol. 104, no. 12: 639-659.

Williamson, Timothy. 2013. 'Gettier Cases in Epistemic Logic.' *Inquiry* vol. 56, no. 1: 1-14.

———. 2017. 'Model-building in Philosophy.' In *Philosophy's Future: The Problem of Philosophical Progress*, eds., Russell Blackford and Damien Broderick. Wiley-Blackwell, 159-173.