Philosophical Studies
https://doi.org/10.1007/s11098-026-02615-w

# No determinate facts, no objective *oughts*?

**Alessandra Marra**[1] · **Christine Tiefensee**[2]

Received: 15 November 2025 / Accepted: 14 September 2026
© The Author(s) 2026

**Abstract**
According to objectivists, what we objectively ought to do depends on the facts. According to defenders of metaphysical indeterminacy, there are cases in which no determinate facts exist. What, then, ought we objectively to do when there are no determinate facts? This paper pursues two objectives in response to this question. The first is constructive: on behalf of the objectivist, we present two proposals for addressing what we call the 'problem from metaphysical indeterminacy'. The second is critical: we argue that, despite their intuitive appeal and close alignment with prominent positions in the literature, neither proposal succeeds. We conclude that objectivists still owe us an account of what we objectively ought to do when no determinate facts obtain.

**Keywords** Objective ought · Subjective ought · Metaphysical indeterminacy · Thin red line · Probabilism

## 1 Introduction

According to objectivists, what we objectively ought to do depends on the facts. According to defenders of metaphysical indeterminacy, there are cases in which no determinate facts exist. What, then, ought we objectively to do in such cases?

Alessandra Marra and Christine Tiefensee contributed equally to this work.

✉ Christine Tiefensee
c.tiefensee@fs.de

Alessandra Marra
alessandra.marra@lrz.uni-muenchen.de

[1] Munich Center for Mathematical Philosophy, Ludwig-Maximilians-Universität München, Munich, Germany

[2] Computational Science & Philosophy Department, Frankfurt School of Finance & Management, Frankfurt, Germany

Published online: 28 October 2026

This paper pursues two goals in response to this question. The first is constructive: on behalf of the objectivist, we present two proposals for addressing what we call the 'problem from metaphysical indeterminacy'. The second is critical: we argue that, despite their intuitive appeal and close alignment with prominent positions in the literature, neither proposal succeeds. We conclude that objectivists still owe us an account of what we objectively ought to do in cases where no determinate facts obtain.

The paper proceeds as follows. Section 2 outlines the success criteria for evaluating potential solutions to the problem from metaphysical indeterminacy by articulating the 'job description' that defines the role of the objective *ought* in normative reasoning. Sections 3 and 4 then present two candidate strategies for addressing this problem. The first draws on the so-called 'thin red line' account from the debate over future contingents to reconcile the objective *ought* with metaphysical indeterminacy. The second proposes a probabilistic approach, suggesting that in cases of metaphysical indeterminacy, what we objectively ought to do is determined not by an action's actual value but by its objectively expected value. We argue that both proposals ultimately fall short. Section 5 concludes with a brief summary of our findings.

Before we begin, let us emphasise that our concern lies exclusively with cases of *metaphysical* indeterminacy, and more specifically with cases in which it is presently indeterminate how the future will unfold. Think, for instance, of ordinary decisions made under an open-future conception of time. Whether one leaves home five minutes earlier than planned, accepts a job offer, or attends a social gathering, there may be no fact of the matter—prior to our decision and the subsequent course of events—which of several possible future histories will obtain as a consequence of our choice. Or consider a government's decision to impose a lockdown during a pandemic. On an indeterministic picture, it may not merely be *unknown* whether the intervention would avert a large outbreak or prove largely ineffective; rather, the future *itself* remains genuinely unsettled. More mundanely, consider a simple bet on a coin toss—the example that will serve as our running illustration throughout this paper. Suppose that, prior to the toss, there is no fact of the matter whether the coin will land heads or tails. It follows that the consequences of the available betting options—whether we will win or lose—are likewise presently indeterminate, even though they will eventually become settled.[1] In all these cases, the issue is not merely epistemic. Even a complete specification of the present may leave multiple future continuations metaphysically open.

We will, therefore, not consider scenarios involving *epistemic* indeterminacy, where for principled epistemic reasons, we lack knowledge of which determinate facts obtain. Nor will we address cases of *semantic* indeterminacy, where concepts lack determinate extensions and would require precisification to resolve indeterminacy in their application. Instead, our focus will be solely on cases of metaphysical indeterminacy, where "independently of how we know or represent the world, there's

[1] Hence, our examples involve coin tosses described at a macroscopical level, in terms of the open metaphysical possibility or the objective chances, at a given time $t_1$, of the coin landing heads or tails at a later time $t_2$. For a picture of metaphysical indeterminacy similar to the one adopted here, and for a defence of the view that metaphysical indeterminacy can exist at the macroscopical level even if metaphysical determinism holds at a microscopical level of description, see List/Pivato (2015).

no determinate way that the world *itself* is" (Williams 2008: p. 763). Whether or not such cases exist remains a matter of controversy. In what follows, we will simply grant, rather than defend, the existence of metaphysical indeterminacy. Given as much, let us ask: What ought we objectively to do when it is metaphysically indeterminate which future outcome will obtain?

## 2 The job description for the objective *ought*

Assessing the plausibility of proposed solutions to the problem from metaphysical indeterminacy requires a prior understanding of the role which the objective *ought* is generally taken to play in normative reasoning. While there is some disagreement over the details, the prevailing view tends to coalesce around the following criteria:

JOB DESCRIPTION

The objective *ought* …

(1) is non-epistemic: What we objectively ought to do depends on what the circumstances of action *are* like, not on what we *believe* them to be like.
(2) corresponds to correctness: The objective *ought* delivers the correct choice. Doing what we objectively ought to do is met with congratulation, whereas failing to do what we objectively ought to do is met with commiseration.
(3) provides third-personal ideal advice: The objective *ought* corresponds to what a fully informed adviser would advise us to do.

The first criterion reflects the distinction between the subjective and the objective *ought*, where the former depends on the information available to the agent, while the latter is insensitive to an agent's epistemic position. "Roughly: what you ought to do in the objective sense has to do with the merits and de-merits of the options available to you, while what you ought to do in the subjective sense has to do with the merits and de-merits of the options available to you, from your epistemic position" (Hare 2011: p. 190).[2]

The second criterion links these two senses of *ought* to corresponding standards of assessment: the subjective *ought* to epistemic appraisals of rationality, and the objective *ought* to non-epistemic standards of correctness. Roughly, "what it is rational to do corresponds to what it makes sense to do given the information that we have, and we are to be blamed for doing what is irrational and praised for its contrary. Whereas what it is correct to do corresponds to what it makes sense to do in light of the facts, … and we are to be congratulated for doing what is correct and commiserated with for doing what is not" (Schroeder 2022: p. 2, see also Schroeder 2018).

The third and final criterion seizes on the widespread thought that the objective *ought* offers third-personal ideal advice. In Judith Thomson's (1986: p. 179) famous

[2] See Carr (2015: p. 678), Gibbard (2005: p. 340), Wedgwood (2003: p. 201; 2016: p. 143), Olson (2017: p. 352).

words: "On those rare occasions on which someone conceives the idea of asking for my advice on a moral matter, I do not take my field work to be limited to a study of what he believes is the case: I take it to be incumbent on me to find out what is the case." Given as much, it is "almost irresistible to think" that the objective *ought* "is the ideal guide, the 'ought' that would guide us if only we knew more, the 'ought' whose advice we most desire" (Muñoz/Spencer 2021: p. 77). If the objective *ought* failed to deliver such advice even under ideal epistemic conditions, it would, as Kolodny and MacFarlane (2010: p. 117) put it, be "useless in deliberation."[3] To avoid misunderstanding, this is not to suggest that ordinary agents can generally appeal to the objective *ought* in first-personal deliberation. Rather, the relevant claim is third-personal: it concerns the advice that an epistemically privileged, fully informed adviser would give.

While these three criteria are closely related, they capture distinct dimensions of the objective *ought*—dimensions that, as we will see later, can come apart. For now, let us conclude that an *ought* qualifies as objective only if it meets all three criteria. Similarly, a solution to the problem from metaphysical indeterminacy will count as objectivist only if it construes *ought* in a manner that is non-epistemic, associated with correctness, and capable of delivering ideal advice.

Next, we present two proposals that aim to satisfy all three desiderata.

## 3 The objective *ought* and the thin red line

Consider the following case and corresponding decision matrix (Table 1):[4]

> *The First Gambling Problem*
> A bookmaker offers us a bet with a €9 stake at $t_1$. Before tossing a coin, she gives us the option either to accept the bet by guessing whether the coin will land heads or tails, or to decline it. If we accept the bet and our guess turns out to be correct at $t_2$, we win €10; if our guess turns out to be incorrect, we lose our €9 stake. At $t_1$, prior to the toss, it is metaphysically indeterminate whether the coin will land heads or tails at $t_2$. The coin is tossed regardless of whether we accept or decline the bet.
>
> What ought we objectively to do?

Fortunately, we are not the first to grapple with cases of metaphysical indeterminacy. Hence, rather than addressing this question from scratch, we can draw on existing debates to see how they might inform what we objectively ought to do when it is metaphysically indeterminate which future outcome will obtain.

One such debate is that concerning future contingents, which examines how to ascribe truth-values to statements about the future if the future is genuinely open. Of

[3] Compare also Wodak (2017: p. 259).

[4] The First Gambling Problem is adapted from Horty (2001: p. 55).

**Table 1** The First Gambling Problem

| | The coin lands tails at $t_2$ | The coin lands heads at $t_2$ |
|---|---|---|
| Bet tails at $t_1$ | 10 euros | 0 euros |
| Bet heads at $t_1$ | 0 euros | 10 euros |
| Decline the bet at $t_1$ | 9 euros | 9 euros |

particular interest here is the so-called 'thin red line' account, which maintains that even in the face of indeterminism, there still is such a thing as the actual future.

To illustrate, consider Fig. 1 and the moment just before the coin is tossed at $t_1$. At that point, it is metaphysically indeterminate on which side the coin will land. Pictorially, let us say that the moment at $t_1$ belongs to several complete histories; these histories share the same past up to $t_1$ but diverge thereafter: in one, the coin lands heads at $t_2$; in another, it lands tails. These are the metaphysically possible futures of $t_1$. However, the thin red line account does not treat these two metaphysically possible futures as equivalent; rather, only one of them represents what will actually happen. The 'thin red line' tracks such actual future. More generally, it marks the complete actual history: it traces what will actually have happened at the end of all times.

As a semantic theory, the thin red line account evaluates future-tense sentences relative to that privileged history. Hence, if the coin lands heads along the history marked by the thin red line, then 'The coin will come up heads' is already true at $t_1$, even though a history on which it lands tails remains metaphysically possible. The appeal of the thin red line account thus lies in its attempt to reconcile two desiderata. On the one hand, it preserves *metaphysical indeterminism* by allowing multiple histories to remain genuinely possible at $t_1$. On the other hand, it preserves the intuition that, among these metaphysically possible futures, exactly one is the *actual future*, so that every future contingent is semantically already either true or false at $t_1$. The account thus combines indeterminism with bivalent truth about what will in fact happen.[5] In the words of Belnap and Green (1994: p. 366):

> Further, belief in an actual future, a Thin Red Line, appears to be consistent with believing that there is indeterminism in the world. There seems nothing untoward in supposing that while there are two things that might become of the coin about to be tossed, it is in fact going to come up Heads. It would seem to be consistent to say, 'The coin will not come up Tails, but it could. What is more, the coin will come up Heads, but it needn't do so'.

Finally, note that the thin red line account does not assume that the actual future is epistemically accessible to us. Rather, for epistemic agents like us, the thin red line is, at best, accessible only in retrospect: only once the coin has landed heads and the

[5] Hence, on thin red line semantics, *truth* and *determinate truth* come apart: 'The coin will land heads' may be true at $t_1$ without being determinately true, that is, without assuming that the coin lands heads along all possible histories passing through $t_1$. By contrast, familiar supervaluationist approaches hold that the sentence 'The coin will land heads' is true at $t_1$ if only if it is determinately true at $t_1$, and false at $t_1$ if and only if it is determinately false at $t_1$. See Thomason (1970) and MacFarlane (2003). Malpass/Waver (2012) propose a novel semantics that combines the thin red line with supervaluationism; however, see Iacona (2014) for criticism.

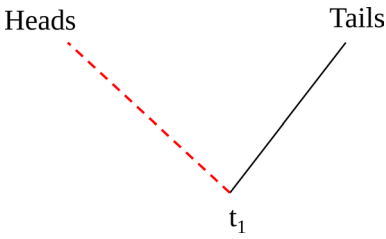


**Fig. 1** The thin red line

outcome has become a determinate fact can we say "Aha! So this is what was going to happen" (Belnap/Green 1994: p. 366; see also MacFarlane 2003: p. 325).

### 3.1 The thin red line *ought*

With the thin red line account in view, let us return to the problem from metaphysical indeterminacy—and with it, the question of what we objectively ought to do when it is metaphysically indeterminate which future outcome will obtain. Borrowing the thin red line framework from the debate over future contingents, here is a first intuitive answer:

> THE THIN RED LINE *OUGHT* (TRL-*OUGHT*)
> In cases of metaphysical indeterminacy, we TRL-ought to perform the action that generates most value in the actual future.[6]

The basic idea is straightforward. If the thin red line account shows that appeal to the actual future enables us to make sense of truth in cases of metaphysical indeterminacy, then perhaps it can also help the objectivist solve the problem posed by metaphysical indeterminacy. More precisely, the proposal is to ground what we objectively ought to do not in presently determinate future facts—after all, there are none—but in the actual future singled out by the thin red line. Applied to the First Gambling Problem: although it is metaphysically open whether the coin will land heads or tails, if the coin will land heads in the actual future, then this is what grounds the verdict that we TRL-ought to bet on heads. Hence, whereas the objective *ought* under metaphysical determinacy is grounded in the determinate facts, the TRL-*ought* under metaphysical indeterminacy is grounded in the actual future.

However, while this idea sounds intuitive, closer inspection shows that the TRL-*ought* does not even get off the ground.[7] To see why, return to the First Gambling Problem and consider its tree-like representation in Fig. 2. To assess this decision situation, the TRL-*ought* offers a maximisation rule: it prescribes that in cases of metaphysical indeterminacy, we TRL-ought to perform the action that generates the most value in the actual future. However, the thin red line account assumes that the actual

[6] The TRL-*ought*—just as the CA-*ought* and the PROB-*ought* discussed below—appeals to the notion of most value. Since our focus lies solely on the problem from metaphysical indeterminacy, we bracket the question of which value system should underpin such an objectivist notion of *ought*. We also set aside issues concerning value ties and value incommensurability.

[7] We thank an anonymous reviewer for their insightful comments and positive suggestions, which prompted the discussion that follows.

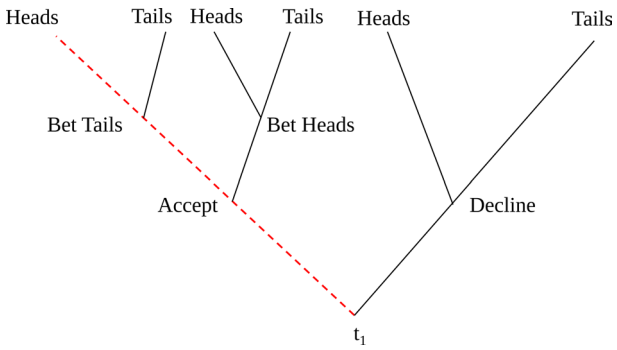


**Fig. 2** Tree-like representation of the First Gambling Problem

future lies on a unique, complete history. The history marked by the thin red line is a complete description of what will happen, including the choice that we will end up making. Hence, of the courses of action open to us, only one occurs in the actual future: if we decide to accept the bet and bet on tails, it is that action, and only that action, that forms part of the thin red line. Consequently, the thin red line contains the future payoff generated by that specific action; however, it assigns no payoffs to the non-actual actions. A meaningful maximisation rule, though, presupposes a value for each option; otherwise, it has no alternatives to compare and hence no range of action-values over which the TRL-*ought* can maximise. It follows that if the only action that occurs on the thin red line is betting on tails, then, trivially, we TRL-ought to bet on tails—even if the coin lands heads and the resulting payoff is €0. Despite its intuitive appeal, then, the TRL-*ought* fails: it delivers only a trivial standard for normative evaluation. What we ought to do collapses into what we will actually do.

### 3.2 The counter-actual *ought*

We have seen that the actual future does not provide the kind of comparative information required for a non-trivial assessment of what we ought to do in cases of metaphysical indeterminacy. After all, the actual future tells us only what *will* happen. To determine what we ought to do, though, we must compare what *would* happen if we performed each of the different actions available to us. Could the TRL-*ought* be revised so as to accommodate this counterfactual information?

One natural thought is to replace the TRL-*ought*'s appeal to the actual future with an appeal to counter-actual futures.[8] Rather than considering what *will* happen, the revised account now asks which future *would* become actual if a particular action were performed. Call this future that action's *counter-actual future*. This yields the following revised principle:

[8] It has long been noted in the literature on future contingents that counterfactuals pose a challenge for thin red line semantics. See Belnap/Green (1994), MacFarlane (2014), Malpass/Wawer (2012), Iacona (2014), and Braüner (2023), who also discuss supplementing thin red line semantics with counter-actual futures.

THE COUNTER-ACTUAL *OUGHT* (CA-*OUGHT*)
In cases of metaphysical indeterminacy, we CA-ought to perform the action whose counter-actual future contains more value than the counter-actual future of every alternative action.

More explicitly, in cases of metaphysical indeterminacy, we CA-ought to perform an action if and only if the future that would become actual were that action performed contains more value than the future that would become actual were any alternative action performed.

Applied to the tree-like representation of the First Gambling Problem, the CA-*ought* thus requires comparing three action-relative futures: the counter-actual future of declining the bet, the counter-actual future of betting on heads, and the counter-actual future of betting on tails.

How should we identify and assign payoffs to these counter-actual futures against the background of metaphysical indeterminacy? To give the CA-*ought* its strongest possible defence, let us grant that the actual future determines not only the payoff of the action that becomes actual, but also the payoffs of the counter-actual futures that correspond to each non-actual action.

More specifically, suppose that each counter-actual future is factually similar to the actual future, in that every element of the actual future is held fixed across counter-actual futures except for the action under consideration. Accordingly, if the coin lands heads in the actual future, then it also lands heads in each counter-actual future.[9]

Applied to the First Gambling Problem and assuming that the leftmost branch of the tree lies on the thin red line, the counter-actual future of betting on tails therefore yields €0, that of betting on heads yields €10, and that of declining the bet yields €9 (see Fig. 3). Hence, although it is metaphysically indeterminate at $t_1$ how the coin will land at $t_2$, the CA-*ought* prescribes that we CA-ought to bet on heads at $t_1$. The CA-*ought* thus provides a non-trivial account of the objective *ought* in situations of metaphysical indeterminacy. The remaining question is: can it also satisfy the job description of the objective ought?

With respect to criterion 1, which demands that the objective *ought* be non-epistemic, and criterion 2, which links the objective *ought* to standards of correctness and their related reactions of congratulation and commiseration, the CA-*ought* shows great promise. Since it grounds its prescriptions *not* in our *beliefs* about what would happen but in what would *actually* happen if an action were performed, the CA-*ought* qualifies as non-epistemic. Moreover, the CA-*ought* also elicits the kinds of reactions typically associated with objective standards of correctness rather than subjective standards of rationality. If, in the First Gambling Problem, we incorrectly decide to bet on tails, our decision will invite commiseration, not blame. Conversely, if we correctly decide to bet on heads, the fitting response will be congratulation, not praise.

[9] Note that, even in simple cases such as the First Gambling Problem, this is a substantive assumption that requires argument rather than stipulation. Moreover, its compatibility with metaphysical indeterminism would itself require further defence (for discussion of the debate over counter-actual futures, see footnote 8). Still, we simply grant this assumption here, as it is especially congenial to the CA-*ought* and thus strengthens its credentials as a solution to the First Gambling Problem.

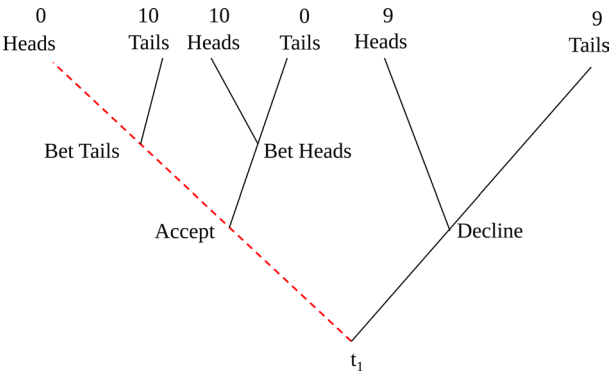


**Fig. 3** Tree-like representation of the First Gambling Problem with payoffs

How about criterion 3, which requires that the objective *ought* offer ideal advice? Let us first get a clearer grasp of what this advisory function involves. Adapting Caspar Hare's (2011: p. 190) suggestions, we can distinguish between two interpretations:[10]

***The Advice of Omniscient Desire*** The CA-*ought* offers ideal advice in the sense that an omniscient, rational being with appropriate interests would want you to do what you CA-ought to do.

***The Advice of Most Reason*** The CA-*ought* offers ideal advice in the sense that what you CA-ought to do is what there is most reason to do.

Can the CA-*ought* offer ideal advice under either of these two interpretations?

The answer is no with regard to the Advice of Most Reason. For the CA-*ought* to count as advisory on this reason-based interpretation, there would need to be most reason at $t_1$ for us to do what the CA-*ought* prescribes. However, there is no such reason. To see why, note that determining what there is most reason to do would require comparing, at $t_1$, the outcomes of the following betting options (see Hare 2011: p. 197):

(A) If you decline the bet, you will retain your €9.
(B) If you accept the bet, it is indeterminate whether the coin will land heads or tails, such that there is no fact of the matter whether you will lose €9 or win €10.

Yet, given the indeterminacy involved in (B), there is no available way to compare these outcomes in a manner that is guaranteed to correspond with the predictions of the CA-*ought*. (We return to the possibility of a probabilistic evaluation of (A) and (B) in the next section.) To elaborate, suppose that we attempt to assess (A) and (B) at $t_1$ using qualitative decision rules from the literature on decision-making under uncertainty, such as Maximin, Maximax, or Minimax Regret (Resnik 1987). Then it

[10] Whereas Hare grounds the objective *ought* itself in omniscient desire and most reason, we associate these considerations only with its advisory function.

quickly becomes apparent that none of these rules ensures alignment with the prescriptions of the CA-*ought*. For instance, Maximin, which recommends choosing the action with the highest minimum value, favours declining the bet, since €9 is the highest minimum payoff. Minimax Regret recommends the same course of action, as declining the bet also incurs the lowest maximum regret, namely only €1.[11] Both rules thus yield recommendations that run directly counter to the CA-*ought*, which—given that the coin lands heads on the thin red line, and thus on all the counter-actual futures—prescribes accepting the bet and wagering on heads.

This result generalises: decision-theoretic rules *cannot* track the prescriptions of the CA-*ought* because these rules are structurally indifferent precisely to those metaphysical facts on which the CA-*ought* depends. Rules such as Maximin, Maximax, or Minimax Regret assess options based on the structural distribution of payoffs across possible states, not based on which of those states is actually going to occur. In contrast, the CA-*ought* discriminates between structurally identical actions—such as betting on heads or betting on tails in Table 1—precisely according to which state obtains on the thin red line. Since standard decision rules are blind to this distinction, none of them is guaranteed to align with the prescriptions of the CA-*ought*. The same structural point extends to more sophisticated decision rules developed for decision-making under indeterminacy.[12] It follows that the advisory function of the CA-*ought* cannot be understood in terms of the Advice of Most Reason.

Does the Advice of Omniscient Desire fare better? Yes. As noted above, the Advice of Omniscient Desire traces the advisory function of the objective *ought* to the desires of an omniscient being. Following Hare (2011: p. 194), let us refer to this being as our Fairy Godmother and assume—again so as to make the strongest possible case for the CA-*ought*—that her omniscience affords complete epistemic access to all truths: nothing escapes the Fairy Godmother's knowledge, including truths about the actual

[11] Minimax Regret recommends the action that leads to the lowest maximum regret, where the maximum regret of an action is calculated as the biggest difference between what the agent can get by choosing an action (e.g., betting on heads) and what the agent could have received if they had chosen a different action (e.g., betting on tails). Thus, in the First Gambling Problem, the maximum regret of betting heads is €10, the maximum regret of betting tails is also €10, while declining the bet has a maximum regret of €1.

[12] For instance, consider *dominance act utilitarianism* as developed by Horty (2001: pp. 59–121), according to which, roughly, an agent ought to do what results from *all* non-strictly dominated actions. Since in the First Gambling Problem, all three available actions are non-strictly dominated yet disagree on whether to bet on heads, it does not follow that we ought to bet on heads, contrary to the CA-*ought*. Horty's dominance act utilitarianism can establish only a weaker claim: 'We ought to bet on heads if the coin will land heads'. But this differs importantly from the CA-*ought*: it is a conditional statement that does not depend on how the indeterminacy of the coin toss is ultimately resolved.

Turning to decision rules from the literature on imprecise credences does not help, either. Let us set aside the issue that these rules may presuppose a probabilistic structure that is not supplied by the thin red line account of metaphysical indeterminacy. The central problem is that, if they are to always match the prescriptions of the CA-*ought*, these rules must guarantee that what we ought to do is the correct action *given what will actually happen*. Adapted to fit the First Gambling Problem, Levi's *E-admissibility* (see Elson forthcoming) yields only the verdict that betting on heads is permissible; Rinard (2015) *Moderate* renders it indeterminate whether betting on heads is permissible; and Williams' (2014) *Randomize* could only occasionally, but not always, coincide with the prescriptions of the CA-*ought*. Hence, none of these rules is guaranteed to be aligned with the CA-*ought*.

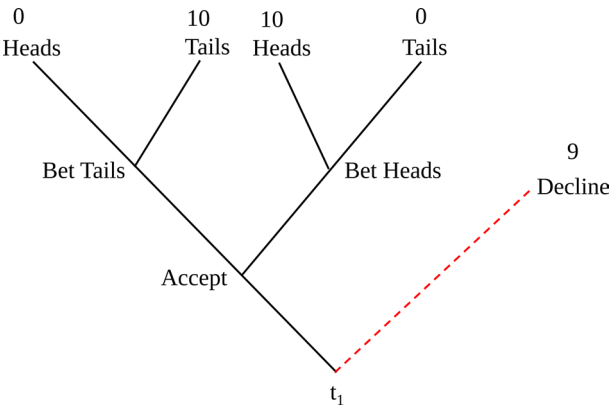


**Fig. 4** Tree-like representation of the Modified First Gambling Problem

future and the counter-actual futures.[13] Hence, although it is metaphysically indeterminate at $t_1$ how the coin will land, and although agents like us lack epistemic access to the future, our Fairy Godmother already knows at $t_1$ what the actual outcome of the coin toss will be. Moreover, assuming factual similarity between the actual future and the counter-actual futures, she thereby also knows what would happen under each relevant counterfactual alternative. Her knowledge therefore includes all of the following truths: that the coin will land heads; that, if you were to bet on heads, you would win €10, that, if you were to bet on tails, you would have €0, and that, if you were to decline the bet, you would have €9.

But if so, then based on this knowledge, the Fairy Godmother can compare what would happen if we were to perform each of the different courses of action available to us, form rational preferences over the respective outcomes, and, in the case at hand, come to desire that we bet on heads at $t_1$. Accordingly, the Advice of Omniscient Desire succeeds where the Advice of Most Reason failed: Since the Fairy Godmother can form the corresponding desire, it secures the advisory function needed for the CA-*ought* to fulfil its job description.

This sounds like excellent news for the objectivist. Yet again, things are not quite as they first seem. Here is why.

So far, we have assumed factual similarity between the actual future and the counter-actual futures, which enabled us to identify the counter-actual future and the corresponding payoff for each non-actual action in the First Gambling Problem. We have already noted that this assumption is substantive. However, even if we continue to grant this assumption, the CA-*ought* still fails to deliver a standard of evaluation that would be widely applicable. To see why, consider a simple modification to the First Gambling Problem:

[13] This characterisation of omniscience in terms of knowledge of all truths aligns with much of the literature on divine omniscience. See, for example, Todd and Rabern (2021), although they proceed from the assumption that future contingents are not true, such that they ultimately do not endorse the thin red line approach.

*The Modified First Gambling Problem*
A bookmaker offers us a bet with a €9 stake at $t_1$. Before tossing a coin, she gives us the option either to accept the bet by guessing whether the coin will land heads or tails, or to decline it. If we accept the bet and our guess turns out to be correct at $t_2$, we win €10; if our guess turns out to be incorrect, we lose our €9 stake. At $t_1$, prior to the toss, it is metaphysically indeterminate whether the coin will land heads or tails at $t_2$. The coin is tossed only if we accept the bet.

In this variant of the First Gambling Problem, everything is as in the original problem except that the coin is tossed *only if* we accept the bet. Suppose further, as represented in Fig. 4, that on the thin red line we in fact *decline* the bet. In that case, there is no actual future in which the coin is tossed, and hence no truth about how the coin will land. But if so, there is also no actual truth that could determine whether the coin would have come up heads or tails had we accepted the bet. As a result, the thin red line can no longer be used to single out the relevant counter-actual futures and corresponding payoffs: it simply does not contain the information required to identify the payoffs needed to apply the CA-*ought*'s maximisation rule. But if so, the CA-*ought* once more reaches a dead end. Even if it is true at $t_1$ what the actual future will hold, and even if the Omniscient Fairy Godmother knows everything there is to know, the CA-*ought* cannot determine what we ought to do, nor can the Fairy Godmother advise us about which action to take.

This leaves us with the following conclusion. The CA-*ought* offers a solution to the problem from metaphysical indeterminacy in the First Gambling Problem, provided that we understand its advisory function in terms of the Advice of Omniscient Desire and that two conditions hold. First, truths about the actual and counter-actual futures must be accessible to an epistemically privileged being. Second, the thin red line must contain the information required to identify the relevant counter-actual futures and their corresponding payoffs. If these conditions are fulfilled, all three criteria of the objective *ought*'s job description are satisfied. If they are not, the CA-*ought* fails to deliver non-trivial prescriptions of what we ought to do in cases of metaphysical indeterminacy.

How limiting are these conditions, particularly the second? We believe: very. Recall the examples mentioned from above: An actual future in which the government decides to keep businesses open cannot help identify the payoffs for counter-actual futures in which the country has been locked down. An actual future in which I decide to stay at home cannot help identify the payoff that would have obtained had I gone to the party. An actual future in which you accept a job in the US cannot help identify the payoff of a counter-actual future in which you accept a job in Europe.

More generally, then, the First Gambling Problem is unusual because a certain event—the coin toss—will occur irrespective of our choices. However, most decision situations—and arguably the most important ones—are more akin to the Modified First Gambling Problem, in that different events will unfold depending on which course of action we choose. Hence, at best, the CA-*ought* has limited applicability, extending only to cases like the First Gambling Problem, in which the actual future contains information about the counter-actual futures. At worst, the very assumption that helped get the CA-*ought* off the ground in the first place is precisely what

undermines its applicability to the most important and prevalent choices we face.[14] The objectivist, then, does not have a widely applicable answer to the problem from metaphysical indeterminacy.

## 4 The objective *ought* and objective probabilities

Even though the TRL-*ought* and the CA-*ought* do not succeed in providing a general solution to the problem from metaphysical indeterminacy, another construal of the objective *ought* might. To introduce this second proposal, consider the following case and corresponding decision matrix (see Table 2):

> *The Second Gambling Problem*
> A bookmaker offers us a bet with a €9 stake at $t_1$. Before she tosses a coin, we can either accept the bet and guess on which side the coin will land, or we can decline the bet. If we accept the bet and our guess turns out to be correct at $t_2$, we win €10. If we accept the bet and our guess turns out to be incorrect at $t_2$, we lose our stake. At $t_1$, it is metaphysically indeterminate whether the coin will land heads or tails. The objective probabilities of it landing heads or tails are 0.5 respectively. The coin is tossed no matter whether we accept or decline the bet.

What ought we objectively to do?
Here is an intuitively plausible answer: In cases of metaphysical indeterminacy, what we objectively ought to do is what is objectively most likely to yield the greatest value. More precisely:

> THE PROBABILISTIC *OUGHT* (PROB-*OUGHT*)
> In cases of metaphysical indeterminacy, we PROB-ought to perform the action that maximises objectively expected value.

Applied to the Second Gambling Problem, the probabilistic construal of the objective *ought* yields a clear result: Since declining the bet maximises objectively expected value, that is precisely what we PROB-ought to do.

Turning to probabilism in the face of metaphysical indeterminacy is widespread, particularly within the consequentialism debate. In this vein, Graham Oddie and Peter Milne (1991: pp. 54–55) declare that in cases of indeterminacy, "the act simply endows various outcomes with various probabilities, and its value must be some kind

[14] The issue is reminiscent of debates about how the antecedents of counterfactual conditionals determine the possible world(s) in which the consequent is to be evaluated (see Stalnaker's seminal 1980). Crucially, in the context of the problem from metaphysical indeterminacy, the choice of a selection mechanism for counter-actual futures bears both on the applicability of the CA-*ought* and on whether it qualifies as the objective *ought*. The objectivist could respond to our argument by dropping the assumption that counter-actual futures are identified on the basis of the information contained in the actual future. While this move is available, it leaves objectivists with a demanding task. They must now provide an alternative selection mechanism that (i) does not trivialise the CA-*ought*, (ii) is not arbitrary, (iii) satisfies the three objectivist criteria, (iv) is compatible with metaphysical indeterminacy, and (v) is applicable both to the First Gambling Problem and to its modified version.

**Table 2** The Second Gambling Problem

| | The coin lands tails at $t_2$, with chance 0.5 | The coin lands heads at $t_2$, with chance 0.5 |
|---|---|---|
| Bet tails at $t_1$ | 10 euros | 0 euros |
| Bet heads at $t_1$ | 0 euros | 10 euros |
| Decline the bet at $t_1$ | 9 euros | 9 euros |

of function of the value of those outcomes and their probabilities given the act", such that "the objective value of an act is just its objective *expected* value." Erik Carlson (1995: pp. 20–21) states that if a certain variant of determinism "is rejected, some form of probabilism would … be more plausible than factualism", while "objective probabilism would … seem preferable to subjective probabilism" as long as there are knowable objective probabilities. Frances Howard-Snyder (1997: p. 241, fn. 1) argues that if "ours is a genuinely indeterministic universe, … then it may be appropriate for objective consequentialism to allow that the right action is the one with the highest 'expected value'", where expected value is determined on grounds of objective probabilities. Arthur Prior (2003: pp. 66–67) explains that "perhaps we can say that if determinism is not true, it suffices to speak of a duty to do what will *probably* have the best total consequences of all the actions open to us", where the probabilities at stake are objective. Jean-Paul Vessel (2007: p. 46) even goes so far as to state that theorists "have consistently maintained that the most plausible forms of objective consequentialism must be probabilistic if and only if indeterminism is true."[15]

The PROB-*ought*, then, enjoys significant popularity. How does it fare with respect to the objective *ought*'s job description? With regard to the first criterion—demanding that the objective *ought* be non-epistemic—things look promising. Unlike subjective accounts, which rely on an agent's credences, the PROB-*ought* appeals to objectively expected value, which is calculated on the basis of objective probabilities. Let us understand these probabilities as single-case chances of event tokens or propositions.[16] For instance, the probability of the coin landing tails refers to the chance, at a given time, that the corresponding outcome occurs—or, equivalently, that the proposition 'The coin lands tails' is true. Accordingly, since the PROB-*ought* rests not on subjective credences but on objective chances, it straightforwardly qualifies as non-epistemic. The first criterion, then, is satisfied.

Things are less straightforward when it comes to criteria 2 and 3. While the PROB-*ought* can offer ideal advice, provided that certain conditions hold, it cannot plausibly

[15] Vessel (2007: p. 46) himself questions the link between indeterminacy and probabilism, arguing that probabilism should be adopted even if determinism is true. Compare Bennett (1998: §§ 17–18) for further thoughts on objective probabilities. See Oddie/Menzies (1992) for a defence of objective probabilistic consequentialism that does not appeal to indeterminacy.

[16] This interpretation can be found in David Lewis (1980), among others, and offers a principled link between the standard objective *ought*, which applies in cases of metaphysical determinacy, and the probabilistic objective *ought*, which applies in cases of metaphysical indeterminacy. On this view, facts grounding the standard *ought* can be understood as event tokens that have obtained. For instance, once the coin has landed tails, the chance of it landing tails becomes 1 and aligns with the truth-value of the corresponding proposition. The standard *ought* thus emerges as an edge case of the probabilistic *ought*, which covers determinate cases by assigning extreme probabilities to the relevant facts. We thank Hannes Leitgeb for discussion.

deliver standards of correctness—or so we shall argue next. We begin with criterion 3.

### 4.1 The PROB-*ought* and ideal advice

Return, then, to matters of ideal advice and the two interpretations we have proposed above. Applied to the PROB-*ought*, they take the following form:

***The Advice of Omniscient Desire*** The PROB-*ought* offers ideal advice in the sense that an omniscient, rational creature with appropriate interests would want you to do what you PROB-ought to do.

***The Advice of Most Reason*** The PROB-*ought* offers ideal advice in the sense that what you PROB-ought to do is what there is most reason to do.

Can the PROB-*ought* offer ideal advice under either of these two interpretations? It can.

Begin with the Advice of Omniscient Desire and return to our Fairy Godmother. On this account, for the PROB-*ought* to fulfil its advisory function, the Fairy Godmother must be able to rationally form desires concerning our betting choice at $t_1$. As before, let us grant that her omniscience extends to all the relevant objective probabilities at $t_1$. Then, given this knowledge, nothing stands in the way of her forming the corresponding desires: as a rational agent, she will desire that we decline the bet, which is precisely what the PROB-*ought* prescribes.

The same positive result holds when the PROB-*ought*'s advisory function is understood in terms of the Advice of Most Reason, albeit with one important qualification. On this interpretation, for the PROB-*ought* to deliver ideal advice, there must be most reason at $t_1$ to do what the PROB-*ought* prescribes. Unlike in the case of CA-*ought*, this condition is met here. For, the Second Gambling Problem now requires us to compare the expected outcomes of the following betting options at $t_1$ (compare Hare 2011: p. 197):

(A) If you decline the bet, then you will retain your €9.
(B) If you accept the bet, it is indeterminate whether the coin will land heads or tails. However, there is a conditional objective probability of 0.5 that you will lose €9 and a conditional objective probability of 0.5 that you will win €10.

This comparison is straightforward: We simply have to calculate and compare the objective expected values of both options. Since this calculation shows that declining the bet maximises objective expected value, we have most reason to decline the bet—which is again exactly what the PROB-*ought* prescribes. Accordingly, the PROB-*ought* can offer ideal advice not only when understood in terms of the Advice of Omniscient Desire, but also when construed along the lines of the Advice of Most Reason.

However, as previously cautioned, this result requires qualification. The reason why is found in a recent argument put forward by Daniel Muñoz and Jack Spencer

(2021) concerning the objective *ought* more generally. To bring this argument to bear on our case, consider a third scenario:

> *The Third Gambling Problem*
> A bookmaker offers us a bet with a €9 stake at $t_1$. Before she tosses a coin, we can either accept the bet and guess on which side the coin will land, or we can decline the bet. If we accept the bet and our guess turns out to be correct at $t_2$, we win €10. If we accept the bet and our guess turns out to be incorrect at $t_2$, we lose our stake. At $t_1$, it is metaphysically indeterminate whether the coin will land heads or tails. At the same time, the coin is heavily biased, with the objective probability of it landing heads being 0.99. While we know that the coin carries a 0.99 bias, we do not know whether it is biased towards heads or tails.

Assume further that the objective and subjective *ought*s are structurally isomorphic, in that the objective PROB-*ought* prescribes maximising objectively expected value, while the subjective *ought* (SUB-*ought*) directs us to maximise subjectively expected value. Applied to the Third Gambling Problem, this entails that we PROB-ought to bet on heads, given that doing so maximises objectively expected value. Knowing that the coin is heavily biased, we also know that we PROB-ought to bet on whichever side the bias favours. It follows that we know that we PROB-ought to accept the bet.[17] Yet, since we lack information about the direction of this bias, our best subjective estimate remains that the coin is equally likely to be biased towards heads or tails. As a result, we SUB-ought to decline the bet, since this option maximises subjectively expected value. Despite the structural isomorphism between the objective and subjective *ought*s, and despite knowing that we PROB-ought to accept the bet, it does not follow that we SUB-ought to do so. Knowing what we PROB-ought to accept the bet, then, does not deliver sufficient advice on how to act.

Importantly, this does *not* show that the PROB-*ought* fails in its advisory function when interpreted in terms of the Advice of Most Reason. Rather, it reveals that this advisory function depends on a specific condition, namely: Epistemically privileged circumstances must include not only *general* knowledge of objective probabilities (e.g., that the coin is heavily biased), but also sufficiently *specific* knowledge of the relevant values (e.g., that the bias is 0.99 towards heads). Absent sufficient specific-

[17] Muñoz/Spencer's (2021: p. 80) own argument draws on what they call the "objective ought's upward monotonicity"—namely, if you ought to do X, and X-ing entails Y-ing, then you ought to do Y—and knowledge closure—namely, if you know that A, and A entails B, then you know that B. Both principles are controversial. For one, the adequacy of knowledge closure has been challenged in the case of non-ideal, bounded epistemic agents, for whom the entailment from A to B may be too complex to grasp (see, e.g., Jago (2006). However, such concerns do not apply in the specific case of the Third Gambling Problem: even for bounded agents, it is plausible that if they know that they ought to bet on some particular outcome, then they must also know that they ought to accept the bet.The principle of *ought*'s upward monotonicity, in turn, is not even required for our purposes. Instead, we can model the scenario as a multi-step decision tree, where the agent first chooses whether to bet or decline and, if betting, then chooses between heads and tails. Reasoning backwards by backward induction—assuming at each step that the agent acts to maximise objectively expected value—we reach the initial decision point with the judgment that betting maximises expected value. Two conclusions follow: Firstly, Muñoz/Spencer's (2021) argument can be developed using less controversial premises. Secondly, their argument applies not only to non-isomorphic objective and subjective *ought*s, but also to the structurally isomorphic PROB-*ought* and SUB-*ought*.

ity, the PROB-*ought* may be knowable in general terms and yet fail to deliver actionable advice.

Despite this qualification, though, our positive result stands: the PROB-*ought*'s advisory function can be captured both by the Advice of Omniscient Desire and the Advice of Most Reason. Hence, whichever interpretation we adopt, criterion 3 of the objective *ought*'s job description is fulfilled. As we will see next, though, the same cannot be said for criterion 2.

### 4.2 The PROB-*ought* and correctness

When introducing the second criterion of the objective *ought*'s job description, we explained that the subjective *ought* is associated with epistemic appraisals of rationality, whereas the objective *ought* is linked with non-epistemic standards of correctness. We also clarified that appraisals of rationality are related to reactions of praise and blame, whereas assessments of correctness invite reactions of congratulation and commiseration. Accordingly, for the PROB-*ought* to meet the second criterion, it must trigger correctness-related reactions of congratulation and commiseration. Does it?

In some cases, yes. Compare:

> *The Fourth Gambling Problem*
> A bookmaker offers us a bet with a €9 stake at $t_1$. Before she tosses a coin, we can either accept the bet and guess on which side the coin will land, or we can decline the bet. If we accept the bet and our guess turns out to be correct at $t_2$, we win €10. If we accept the bet and our guess turns out to be incorrect at $t_2$, we lose our stake. At $t_1$, it is metaphysically indeterminate whether the coin will land heads or tails. At the same time, the coin is heavily biased, with the objective probability of it landing heads being 0.99. While we know that the coin carries a 0.99 bias, we do not know whether it is biased towards heads or tails. The coin is tossed and lands heads at $t_2$.

As we argued above, in this case all of the following are true: We SUB-ought to decline the bet; we PROB-ought to bet on heads; we know that we PROB-ought to accept the bet; and yet we do not know whether we PROB-ought to bet on heads or on tails.

Assume that we decide to bet on heads—in line with the prescription of the PROB-*ought*, but against the verdict of the SUB-*ought*. Then although we now deserve blame for having chosen irrationally, we should nevertheless be congratulated at $t_2$ for having made the correct betting decision. So far, so good.
However, now consider the following case:

> *The Fifth Gambling Problem*
> A bookmaker offers us a bet with a €9 stake at $t_1$. Before she tosses a coin, we can either accept the bet and guess on which side the coin will land, or we can decline the bet. If we accept the bet and our guess turns out to be correct, we win €10. If we accept the bet and our guess turns out to be incorrect, we lose our

stake. If we refrain from betting, we keep our original stake of €9. Although we believe that the coin is 0.99 biased towards tails, it is in fact biased 0.99 towards heads. The coin is tossed and lands tails at $t_2$.

In this case, we SUB-ought to bet on tails but PROB-ought to bet on heads at $t_1$. At $t_2$, though, we will not only deserve praise for having chosen rationally, but also be *congratulated* for betting on tails—despite the fact that we PROB-ought to have bet on heads!

Indeed, we can go further. Consider:

*The Sixth Gambling Problem*
A bookmaker offers us a bet with a €9 stake at $t_1$. Before she tosses a coin, we can either accept the bet and guess on which side the coin will land, or we can decline the bet. If we accept the bet and our guess turns out to be correct, we win €10. If we accept the bet and our guess turns out to be incorrect, we lose our stake. If, in turn, we refrain from betting, we keep our original stake of €9. We correctly believe that the coin carries a 0.99 bias towards landing tails. The coin is tossed and lands heads at $t_2$.

In this case, we both SUB-ought and PROB-ought to bet on tails. Yet, despite having chosen rationally and having followed the prescriptions of both the SUB-*ought* and the PROB-*ought*, we will now face *commiseration* at $t_2$!

These discrepancies between the prescriptions of the PROB-*ought* and correctness-based reactions of congratulation and commiseration are no accident. Rather, they arise from the simple fact that even objectively expected outcomes can diverge from actual outcomes. Given this divergence, acting in line with what we PROB-ought to do can turn out to be the *incorrect* choice, just as acting against the PROB-*ought* can turn out to be the *correct* choice. But if so, it is no longer clear what role the PROB-*ought* is supposed to play within our normative practices. Since it is neither guaranteed to align with actual outcomes at $t_2$ nor with what we SUB-ought to do at $t_1$, it neither addresses objective assessments of correctness nor subjective assessments of rationality. What normative role, then, is left for the PROB-*ought*?

Here are two possible replies.[18] The first preserves the PROB-*ought*'s traditional correctness function and argues that our considerations leave this function intact. The second concedes that our considerations show that the PROB-*ought* does not fulfil its traditional correctness function and assigns it a different normative role instead.

Begin with the first reply. When discussing the Gambling Problems above, we argued that the PROB-*ought* does not qualify as the objective *ought* because it fails what we can call the 'reaction test': since objectively expected outcomes can diverge from actual outcomes, acting in accordance with the PROB-*ought* need not elicit the reactions of congratulation and commiseration that are indicative of correct and incorrect choices. The first reply now maintains that this argument applies the reaction test to the wrong object. Our discussion focused on the reactions elicited by

[18] We are grateful to an anonymous reviewer for suggesting a response along the lines of the second reply discussed here.

the realised outcome of the gamble. However, the PROB-*ought* identifies correctness with the maximisation of objectively expected value. Accordingly, reactions of congratulation and commiseration should track the latter, not the former. An action invites congratulation if it maximises objectively expected value; it invites commiseration if it does not. The actual outcome of the coin toss is neither here nor there. Hence, the PROB-*ought does* pass the reaction test and thus qualifies as the objective *ought*. Our argument merely applied this test to the wrong object, or so critics now argue.

While available, we believe that this reply is costly. Much of the initial appeal of the objective *ought* stems precisely from the intuition that an action which is expected to maximise value can nevertheless turn out to be wrong. By identifying correctness with the maximisation of expected value, even if this is understood as objectively expected value, the first reply fails to capture this intuition. As a result, it undermines one of the key motivations for introducing the objective *ought* in the first place.

The second reply agrees that correctness should continue to track actual value rather than objectively expected value. It therefore also concedes that the PROB-*ought* does not fulfil the traditional correctness function of the objective *ought*. Instead, it suggests that the PROB-*ought* serves a different normative role: it does not answer the question 'What ought we to do?', but the different question 'What ought we to do if we know everything there is to know?'. Since this question concerns the actions of an ideally informed agent, rather than those of agents with limited information, it retains an important sense of objectivity. And if, in cases of metaphysical indeterminacy, all there is to know consists in the relevant objective probabilities, the PROB-*ought* qualifies as objective precisely because it identifies what an ideally informed agent ought to do in light of that knowledge. Acting in accordance with the PROB-*ought* thus amounts to doing what we ought to have done had we possessed complete information.

This reply would need to be fleshed out further. For instance, does the PROB-*ought*, thus understood, function as an ideal guide to the SUB-*ought* (bearing our earlier remarks about the specificity of information in mind)? Does acting in accordance with its prescriptions merely shield us from blame—a reaction traditionally associated with rationality—or is some kind of correctness in play after all? However, even if we grant that the normative role of the PROB-*ought* can be further sharpened, our initial concern remains. We wanted to know what we objectively ought to do in cases of metaphysical indeterminacy. The second reply instead tells us what we ought to do in such cases if we were fully informed.

Hence, although probabilism is widely regarded as the leading response to metaphysical indeterminacy, and although the PROB-*ought* is both non-epistemic and capable of offering ideal advice, we face the following dilemma. Either, objectivists build probabilities into the very notion of objective correctness. If so, they forgo one of the central motivations for introducing the objective *ought* in the first place. Or they concede that the PROB-*ought* does not fulfil the objective *ought*'s traditional correctness function and instead assign it a different normative role. But if so, the PROB-*ought* no longer answers the original question of what we objectively ought to do in cases of metaphysical indeterminacy. Appealing to the PROB-*ought*, then,

either offers a costly solution to the problem from metaphysical indeterminacy or changes the subject.

### 4.3 Correctness and time-point stability: the special case of the PROB-*ought*

At this point, though, followers of the debate about the objective *ought* might find themselves with a nagging question. For, as Mark Schroeder (2018, 2022) has recently argued, even in contexts of metaphysical determinacy, the objective *ought*'s connection to standards of correctness—and to the correctness-related reactions of congratulation and commiseration—might be less stable than commonly assumed.

According to Schroeder, objective *ought* judgements legitimately differ depending on the informational perspective from which they are made. An adviser with partial information may therefore judge that an agent objectively ought to perform one action, while a better-informed adviser may judge that she objectively ought to perform another. The objective *ought* thus no longer functions as a single, information-invariant standard of correctness. Still, it remains distinct from rationality, which continues to evaluate an agent's decisions based on the agent's own information.

Let us grant—if only for the sake of argument—that Schroeder is right in his general diagnosis of the objective *ought*. Does it follow that the concerns about correctness raised in the previous section are not distinctive of the PROB-*ought*, but instead reflect a more general feature of the objective *ought*?

We do not think so. For, even if we grant that Schroeder is right about the objective *ought*, we will argue that the problems afflicting the PROB-*ought* are importantly different—and more troubling—than Schroeder's broader observations. Accordingly, our assessment of the PROB-*ought* remains unaffected by recent developments in the debate about the objective *ought*.

To elaborate, consider Schroeder's 'four-envelope case', which pertains to the objective *ought* in contexts of metaphysical determinacy:

> *The Four-Envelope Case (Schroeder*2018, 2022)
> Xiao is tasked with choosing one of four envelopes, which contain \$0, \$4000, \$5000 and \$6000 respectively. While it is common knowledge that each envelope contains one of these amounts, Xiao, Ying, and Zach differ in how much information they have about the envelopes' contents:
> Xiao knows at $t_1$ that envelope 1 contains \$4000. From Xiao's perspective, she ought to choose envelope 1.
> Ying knows at $t_1$ that envelope 1 contains \$4000 and that envelope 2 contains \$5000. From Ying's perspective, Xiao ought to pick envelope 2.
> Zach knows at $t_1$ the exact contents of each envelope: envelope 1 contains \$4000, envelope 2 \$5000, envelope 3 \$6000, and envelope 4 \$0. From Zach's perspective, Xiao ought to pick envelope 3.
> Xiao picks the envelope at $t_1$. The contents of the envelopes are revealed at $t_2$.

Let us focus on Ying in particular, who, based on her knowledge at $t_1$, advises Xiao to take envelope 2. Granting Schroeder's suggestion that Ying's *ought*-judgements

count, in *some* relevant sense, as objective, let us say that Xiao YING-ought to take envelope 2.

Now suppose that Xiao instead picks envelope 3, which at $t_2$ is revealed to contain \$6000. Then at $t_2$, Ying should *congratulate* Xiao on having picked envelope 3, even though Xiao *failed* to do what she YING-ought to have done. Contrary to expectations about correctness-related reactions, then, violations of the YING-*ought* can warrant congratulation rather than commiseration—just as we observed in the case of the PROB-ought!

However, even if we grant this analogy between the YING-*ought* and the PROB-*ought*, an important disanalogy remains. To see why, we need to distinguish between two different kinds of objective *ought*-judgements and two diverging senses of correctness.

Starting with the former, note that at $t_1$, Ying not only judges that Xiao YING-ought to take envelope 2, but also that she YING-ought not to take envelope 1. Now suppose that—while all the facts remain fixed—a trustworthy source whispers to Ying that envelope 3 contains more money than envelope 2, thereby updating her informational state. Then *at* $t_1$, Ying would now revise her judgement that Xiao YING-ought to pick envelope 2, but would not revise her judgement that she YING-ought not to take envelope 1. Let us call objective *ought*-judgements that are sensitive to such informational updates '*time-point unstable*' and those that remain fixed in light of new knowledge '*time-point stable*'.

Similarly, let us distinguish between two senses of correctness. Let us say that an objective *ought*-judgement secures *strong correctness* if acting in accordance with it guarantees making the correct choice, whereas an objective *ought*-judgement secures *weak correctness* if acting against it guarantees making an incorrect choice.

Combining both distinctions, we can now see that, in the case of the general objective *ought*, time-point unstable objective *ought*-judgements are neither strongly nor weakly correct: following Ying's judgement that Xiao YING-ought to take envelope 2 neither guarantees that Xiao will make the correct choice, nor does acting against it guarantee that she will make an incorrect choice. By contrast, time-point stable objective *ought*-judgements are at least weakly correct: failing to follow Ying's judgement that Xiao YING-ought not to take envelope 1 will always constitute an incorrect choice, regardless of any further information Ying might acquire about the contents of envelopes 3 and 4.

Matters stand differently with the PROB-*ought*. First, note that all PROB-*ought* judgements are time-point stable, given that actual outcomes do not affect objective probabilities: If the objective chance of the coin landing heads is 0.5, and the coin does in fact land heads, this does not change the fact that, *at* $t_1$, the objective chance of heads was still only 0.5. More generally, since the PROB-*ought* is grounded in objective chances assigned to event tokens, and since the actual outcome of the coin

toss does not impact these objective chances *at* $t_1$, learning about this outcome does not lead to a (retrospective) revision of the original PROB-ought judgement at $t_1$.[19]

Next, note that given our discussions above, we already know that PROB-*ought* judgements are not strongly correct: As the Sixth Gambling Problem has shown, acting in line with the PROB-*ought* does not ensure choosing correctly. However, observe, finally, that despite being time-point stable, PROB-*ought* judgements are not even weakly correct: As was evident in the Fifth Gambling Problem, failing to act in line with the PROB-*ought* does not ensure choosing incorrectly, either.

The reason lies, once again, in the specific structure of the PROB-*ought*: Since an objectively expected outcome is not guaranteed to occur, acting against the PROB-*ought* does not necessarily result in an incorrect choice, just as acting in accordance with it does not necessarily yield a correct one. Hence, even if we grant, with Schroeder, that the connection between the objective *ought* and correctness is more complicated than assumed, this does not undermine our case against the PROB-*ought*: Whereas time-point stable objective *ought*-judgements are at least weakly correct, time-point stable PROB-*ought* judgements are neither weakly nor strongly correct. Accordingly, our verdict stands: the role of the PROB-*ought* in our normative practices remains obscure. The PROB-*ought* thus fails to offer a viable solution to the problem from metaphysical indeterminacy.

## 5 Conclusion

In this paper, we have argued that objectivists face the problem from metaphysical indeterminacy: they must explain what we objectively ought to do when no determinate facts obtain. We examined two prominent strategies for addressing this problem, one drawing on the thin red line account from debates about future contingents and the other appealing to probabilism. Although both approaches enjoy considerable support and preserve important features of the objective *ought*, neither provides a satisfactory, widely applicable solution. The problem from metaphysical indeterminacy therefore remains unresolved: objectivists still owe us an account of what we objectively ought to do when no determinate facts obtain.

**Acknowledgements** We thank Dominik Klein, Hannes Leitgeb, Sebastian Köhler and two anonymous reviewers for their insightful and constructive comments on previous versions of this paper. AM was supported by the 'Logic of Reasons' project, Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 545459979.

[19] We borrow this observation from Lewis's (1980) 'crystal ball' example, developed in the context of the Principal Principle and the notion of admissible information that may affect the objective chances of event tokens (see also Vranas (2004: p. 375). In the present context, this observation can be supported by means of a dilemma. Suppose that the coin will land heads at $t_2$. Does this information retrospectively alter the objective chances, *at* $t_1$, of the fair coin landing heads? If the answer is yes, then the chance of the coin landing heads becomes 1 at $t_1$, thus reflecting the actual outcome of the coin toss. But in that case, the PROB-*ought* collapses into the TRL-*ought*, thereby inheriting the problems discussed in Sect. 3. If the answer is no, then the PROB-*ought* remains distinct from the TRL-*ought*, yet suffers from the problems we discuss in Sect. 4.

**Funding** Open Access funding enabled and organized by Projekt DEAL.

## Declarations

**Competing interests** There are no competing interests to declare.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

## References

Belnap, N., & Green, M. (1994). Indeterminism and the thin red line. *Philosophical Perspectives,8*, 365–388.
Bennett, J. (1998). *The Act Itself*. Clarendon Press.
Braüner, T. (2023). The true futures. *Synthese,202*(162), 1–23.
Carlson, E. (1995). *Consequentialism Reconsidered*. Springer.
Carr, J. (2015). Subjective Ought, *Ergo* 2: 678–710.
Elson, L. (forthcoming). *Vagueness in Action: An Indeterminist Account of Preference and Choice*. Oxford University Press.
Gibbard, A. (2005). Truth and correct belief. *Philosophical Issues,15*, 338–350.
Hare, C. (2011). 'Obligation and Regret When There is No Fact of the Matter About What Would Have Happened if You Had not Done What You Did', *Noûs* 45: 190–206.
Hare, C. (2020). Consequentialism and indeterminacy. In D. W. Portmore (Ed.), *The oxford handbook of consequentialism* (pp. 344–357). Oxford University Press.
Horty, J. F. (2001). *Agency and Deontic Logic*. Oxford University Press.
Howard-Snyder, F. (1997). The rejection of objective consequentialism. *Utilitas,9*, 241–248.
Iacona, A. (2014). Ockhamism without thin red lines. *Synthese,191*, 2633–2652.
Jago, M. (2006). Hintikka and Cresswell on logical omniscience. *Logic and Logical Philosophy,15*, 325–354.
Kolodny, N., & MacFarlane, J. (2010). Ifs and oughts. *Journal of Philosophy,107*, 115–143.
Lewis, D. (1980). A subjectivist's guide to objective chance. reprinted with postscripts in Lewis, D (1986). *Philosophical Papers: Volume II*. Oxford University Press.
List, C., & Pivato, M. (2015). Emergent chance. *Philosophical Review,124*, 119–152.
MacFarlane, J (2014). *Assessment sensitivity: Relative truth and its applications*. Oxford University Press.
MacFarlane, J. (2003). Future contingents and relative truth. *Philosophical Quarterly,53*, 321–336.
Malpass, A., & Wawer, J. (2012). A future for the thin red line. *Synthese,188*(1), 117–142.
Muñoz, D., & Spencer, J. (2021). Knowledge of objective 'oughts': Monotonicity and the new miners puzzle. *Philosophy and Phenomenological Research,103*, 77–91.
Oddie, G. and Milne, P. (1991). Act and value: Expectation and the representability of moral theories, *Theoria* 57: 42–76.
Oddie, G., & Menzies, P. (1992). An objectivist's guide to subjective value. *Ethics,102*, 512–533.
Olson, K. (2017). A defense of the objective/subjective moral ought distinction. *The Journal of Ethics,21*, 351–373.
Prior, A. N. (2003). *Papers on Time and Tense*. Edited by Per Hasle, Peter Øhrstrøm, Torben Braüner, and Jack Copeland. Oxford University Press.
Resnik, M. D. (1987). *Choices: An Introduction to Decision Theory*. University of Minnesota Press.
Rinard, S. (2015). A decision theory for imprecise probabilities. *Philosophers' Imprint,15*(7), 1–16.

Schroeder, M. (2018). Getting perspective on objective reasons. *Ethics,128*, 289–319.
Schroeder, M. (2022). 'Rationality in Retrospect.' In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. Volume 17, pp. 1–26). Oxford University Press.
Stalnaker, R. (1980). A defense of conditional excluded middle. In: Harper, Stalnaker, & Pearce (eds), *Ifs: Conditionals, belief decision, chance and time*. The Netherlands: Kluwer Academic Publishers
Thomason, R.H. (1970). Indeterminist time and truth-value gaps, *Theoria*, 36: 264–281.
Thomson, J. J. (1986). *Rights, restitution, and risk: Essays in moral theory*. Cambridge: Harvard University Press. Edited by William Parent.
Todd, P., & Rabern, B. (2021). Future contingents and the logic of temporal omniscience. *Noûs,55*, 102–127.
Vessel, J.-P. (2007). The probabilistic nature of objective consequentialism. *Theoria,73*, 46–67.
Vranas, P. B. M. (2004). Have your cake and eat it too: The old principal principle reconciled with the new. *Philosophy and Phenomenological Research,69*, 368–382.
Wedgwood, R. (2003). Choosing Rationally and Choosing Correctly. In Sarah Stroud, S. and Tappolet, C. (eds.) *Weakness of Will and Practical Irrationality*. Oxford: Oxford University Press: 201–229.
Wedgwood, R. (2016). 'Objective and Subjective "Ought."' In N. Charlow & M. Chrisman (Eds.), *Deontic Modality* (pp. 143–168). Oxford University Press.
Williams, R. G. (2008). Ontic vagueness and metaphysical indeterminacy. *Philosophy Compass,3*, 763–788.
Williams, R. G. (2014). Decision-making under indeterminacy. *Philosophers' Imprint,14*(4), 1–34.
Wodak, D. (2017). Can objectivists account for subjective reasons? *Journal of Ethics & Social Philosophy,12*, 259–279.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.