

An Instrumentalist Account of How to Weigh Epistemic and Practical Reasons for Belief

Asbjørn Steglich-Petersen · Mattias Skipper
Forthcoming in *Mind*

Abstract. When one has both epistemic and practical reasons for or against some belief, how do these reasons combine into an all-things-considered reason for or against that belief? The question might seem to presuppose the existence of practical reasons for belief. But we can rid the question of this presupposition. Once we do, a highly general ‘Combinatorial Problem’ emerges. The problem has been thought to be intractable due to certain differences in the combinatorial properties of epistemic and practical reasons. Here we bring good news: if we accept an independently motivated version of epistemic instrumentalism—the view that epistemic reasons are a species of instrumental reasons—we can reduce The Combinatorial Problem to a relatively benign problem of how to weigh different instrumental reasons. As an added benefit, the instrumentalist account can explain the apparent intractability of The Combinatorial Problem in terms of a common tendency to think and talk about epistemic reasons in an ‘elliptical’ manner.

1. Introduction

We often find ourselves in need of weighing different reasons for and against holding a particular belief on some matter. Most obviously, this happens when different parts of our evidence pull in opposite directions with respect to the same proposition. A physician might, for example, encounter opposing studies of whether Pancreatic Inflammation causes Type-2 Diabetes in elderly patients. Here the task is to weigh certain *epistemic* reasons against certain other *epistemic* reasons. Difficult as it may be, the challenge is a familiar one, and one that has received plenty of attention in the epistemological literature.

Sometimes, however, we face a different kind of challenge, namely that of weighing *epistemic* reasons against reasons of a distinctly *practical* kind. Mundane examples are easy to come by: believing that I’m a uniquely gifted tennis player might help me win the next match; believing in the existence of an afterlife might help alleviate grief over the death of a close relative; and so on. But to keep matters relatively simple, let us consider an example inspired by Kavka’s (1983) famous ‘toxin puzzle:’

Dr. Fisch and the Eccentric Billionaire: Dr. Fisch is an experienced fisherman. Based on the mild winter weather, he believes that the herring run will come early this year. This morning, however, Dr. Fisch runs across Bill, an eccentric billionaire who is in the habit of squandering his wealth on erratic wagers. Bill offers Dr. Fisch \$1M if he, in whatever way he sees fit, forms the belief that the herring run will come late this year.

Let us straightaway assume that Bill's offer provides Dr. Fisch with a practical reason to take whatever means he sees fit to form the belief that the herring run will come late this year. That is just to say that Dr. Fisch is not indifferent to the prospect of monetary gain.¹ Let us also assume that Dr. Fisch's belief that the herring run will come early this year is based on good evidence (say, some extensive behavioral studies of herring shoals). Given this, Dr. Fisch seems to face a weighing challenge. If he responds successfully to his practical reasons alone, he will end up believing that the herring run will come late this year. If he responds successfully to his epistemic reasons alone, he *won't* end up with this belief. Thus, the question arises: what does Dr. Fisch have all-things-considered reason to believe?

This question touches upon a long-standing debate in epistemology about whether putative practical reasons for belief are best understood as reasons for belief *per se*, or whether they should be understood in some other way. The contemporary debate on this issue is dominated by two general views. On the one hand, the 'pragmatists' hold that one can indeed have genuine practical reasons for belief. According to them, the normative status of a belief depends not only on its epistemic properties (such as its truth or relation to the evidence), but also depends on its role in the agent's practical life (such as its conduciveness to monetary gain).² On the other hand, the 'evidentialists' hold that all reasons for belief are epistemic in nature. According to them, the normative status of a belief is only sensitive to epistemic factors, and hence is detached from whatever practical costs or benefits might go along with holding that belief.³ Instead, evidentialists typically seek to understand putative practical reasons for belief, not as reasons to *believe* as such, but as reasons to *act* in ways that will 'induce' or 'bring about' the relevant belief. For example, I

¹ It is obviously controversial whether, or to what extent, it is psychologically possible to revise one's beliefs in response to practical considerations. But without getting our hands dirty in this debate, we can assume that Dr. Fisch has *some* means to form the relevant belief. Perhaps he possesses an unusually high degree of voluntary control over his beliefs, or perhaps he knows a remarkably competent hypnotist who is able to manipulate his beliefs.

² Recent proponents of pragmatism include Reisner (2008; 2009; ms.), McCormick (2015), and Rinard (forthcoming).

³ Recent proponents of evidentialism include Adler (2002), Shah (2006), and Berker (2018).

might have practical reason to take a belief-manipulating drug that causes me to believe that the earth is flat.

Note, however, that even if the prospect of monetary gain only provides Dr. Fisch with a practical reason to *bring about* the relevant belief, successfully responding to this reason is still going to preclude him from successfully responding to his epistemic reasons. Thus, Dr. Fisch faces a weighing challenge regardless of whether his epistemic and practical reasons strictly speaking bear on the same thing. So, irrespective of how the debate between pragmatists and evidentialists pans out, both camps face a pressing, yet surprisingly underexplored, question prompted by cases like that of Dr. Fisch and the Eccentric Billionaire:

The Combinatorial Problem: When one possesses both epistemic and putative practical reasons for or against some belief, how do these reasons combine into an all-things-considered reason for or against that belief?

We have deliberately formulated The Combinatorial Problem in a way that leaves open whether the ‘putative practical reasons’ are genuine reasons for belief or not. This should help remind us that the problem is a general one that arises for pragmatists and evidentialists alike. In the end, our own solution to The Combinatorial Problem will not stay entirely neutral in this debate. But the problem itself does not stick its neck out with respect to the divide between evidentialists and pragmatists.

By contrast, The Combinatorial Problem *does* presuppose that epistemic and putative practical reasons for belief combine so as to give rise to an all-things-considered normative verdict in cases like that of Dr. Fisch and the Eccentric Billionaire. Someone might deny this presupposition, and instead maintain that Dr. Fisch faces a kind of normative dilemma, which forces him into either practical or epistemic irrationality. But we suspect that few philosophers would want to dismiss The Combinatorial Problem on this ground.⁴ After all, normative dilemmas are notoriously puzzling, and so it seems worthwhile to investigate whether they can be avoided in the present context. That being said, our aim is not to argue that epistemic and putative practical reasons for belief do indeed combine to yield all-things-considered normative verdicts. The principal aim of this paper is not to defend the presuppositions of The Combinatorial Problem, but rather to investigate how the problem is best solved given that it arises. To the extent that we

⁴ An exception is Feldman (2000, pp. 192-94) who rejects the existence of all-things-considered reasons for belief, or ‘just plain oughts’ as he calls them. See Booth (2012) for a critical reply to Feldman.

succeed, this may be taken to provide some indirect grounds for thinking that Dr. Fisch does not face a genuine normative dilemma.

Recently, some authors have found The Combinatorial Problem to be intractable due to certain differences in the combinatorial properties of epistemic and practical reasons. One such difference becomes salient in situations where one has equally strong, but opposing, reasons of either an epistemic or a practical kind. If one has equally strong *practical* reason to perform each of two incompatible actions (e.g., spending one's money on ice-cream or saving it for later), and no other reasons bearing on the matter, it seems intuitive to say that either course of action is permitted: you may have ice-cream, and you may save your money for later. This combinatorial property is what Selim Berker (2018) calls 'permissive balancing.' By contrast, if one has equally strong *epistemic* reason to believe and disbelieve the same proposition, it does *not* seem intuitive to say that one may believe either way. Rather, one should suspend judgment on the matter. This combinatorial property is what Berker calls 'prohibitive balancing.'

Another notable difference between epistemic and practical reasons concerns the *scales* on which their weights appear. As Andrew Reisner (ms.) has pointed out, we are used to thinking that epistemic reasons can be *conclusive* in the sense that no further evidence can confer any additional support to the relevant belief. In other words, the weights of epistemic reasons seem to feature on a *bounded* scale. By contrast, we are used to thinking that the weights of practical reasons feature on an *unbounded* scale. For example, when it comes to promoting well-being, there is no apparent upper limit beyond which no further strengthening of practical reason can be gained.

One might take these apparent differences to suggest that epistemic and practical reasons are not at all amenable to mutual comparison. Indeed, Berker has contended that epistemic and practical reasons are 'like oil and water: their combinatorial behaviors are just too different for them to mix together in a well-behaved manner' (2018, p. 434). For reasons that will emerge, we do not think that the apparent combinatorial differences show anything this strong. But they do at least suggest that The Combinatorial Problem constitutes a serious challenge worth addressing.

How, then, might we go about solving The Combinatorial Problem? Some authors have tried to define a general rule or procedure for how to combine epistemic and practical reasons. For example, Reisner (2008) has proposed a 'defeating function,' which maps sets of epistemic and practical reasons to verdicts about what there is all-things-considered reason to believe. However, there are serious problems associated with the attempt to specify such a weighing function. Most obviously, the strategy of reverse-engineering a weighing function with certain desirable properties

runs a risk of *ad hocness*.⁵ Even if it turns out possible to encode many (or all) of our intuitions about how epistemic and practical reasons combine in a weighing function, it is doubtful that such a gerrymandered function would be of much theoretical or explanatory value. This is not meant to discourage further attempts to specify a weighing function that is both extensionally adequate and theoretically illuminating. But as things stand, the strategy does not strike us as very promising.

In this paper, we will instead try a different tack on The Combinatorial Problem. Rather than approaching the problem ‘head-on’, we will offer a way of *reducing* the problem to a different problem of a less demanding kind. In short, our claim is this: if we accept an independently motivated version of epistemic instrumentalism—the view that epistemic reasons are a species of instrumental reasons—we can treat The Combinatorial Problem as an instance of a relatively benign problem of how to weigh different instrumental reasons. By thus reducing The Combinatorial Problem, we will not have offered a complete solution to it. After all, it may be challenging to balance conflicting instrumental reasons. But at least this challenge does not seem intractable in the way that The Combinatorial Problem has appeared.

Like any substantive philosophical view, epistemic instrumentalism comes with its own potential follies and foibles. The view has many notable attractions, but we do not take it to constitute anything like a theoretically neutral starting point. In an obvious sense, this is a weakness of our proposed solution, since it won’t be of much use to staunch opponents of epistemic instrumentalism. But at the same time, if it turns out that epistemic instrumentalism can help us address The Combinatorial Problem, this is itself a point in favor of the view.

The rest of the paper proceeds as follows: in §2, we provide some background on epistemic instrumentalism. In §3, we develop our instrumentalist solution to The Combinatorial Problem, and show how the instrumentalist account can help explain the apparent intractability of The Combinatorial Problem in terms of a common tendency to think and talk about epistemic reasons in an ‘elliptical’ manner. In §§4-5, we defend our account against some potential worries. In §6, we conclude.

2. Epistemic Instrumentalism

Why does evidence for a proposition speak in favor of believing it? And why are beliefs in evidentially unsupported propositions to be avoided? These are questions that any theory of

⁵ See Berker (2018) for discussion.

epistemic normativity must answer. Epistemic instrumentalists answer them in terms of familiar ideas about the reasons we have to take means to our aims: we often want, or have reason to want, to believe what is true and avoid believing what is false. And since the best means to this aim is to believe in accordance with the evidence, we have good instrumental reason to do so.

Different instrumentalist accounts differ in both content and motivation. Early accounts were largely motivated by a desire to make epistemic evaluations compatible with a kind of naturalism that is otherwise hostile to normativity. Such accounts grounded epistemic reasons in the actual aims or desires of agents, thereby turning the question of what there is reason to believe into a question of whether certain factual matters obtain; facts about an agent's psychological state and how the satisfaction of that state is best pursued.⁶ However, while such accounts are tidy in their metaphysical commitments, they are widely regarded as inadequate as a general theory of epistemic normativity. The basic problem is that actual epistemic agents don't have enough, or sufficiently general, epistemic aims and desires to generate all of the epistemic reasons they seem to possess.⁷ There are propositions that we are utterly indifferent about (is the number of herrings in the Pacific Ocean odd or even?), and even propositions that we would prefer *not* to form true beliefs about (how does the movie end?). Yet, even in such cases, it seems that our beliefs are normatively constrained by the evidence. The contingency of people's actual aims and desires thus seems to make them unable to explain the generality with which evidence gives rise to normative requirements.

More recent instrumentalist accounts ground epistemic normativity, not in people's *actual* aims, but in the aims they have *reason* to pursue.⁸ This approach gives up on the hope of making epistemic normativity easily adaptable by a naturalist metaphysics, insofar as the question of which aims one has reason to pursue is a normative one. But it still promises a different kind of theoretical unity, namely between practical and epistemic normativity. There are several ways to carry out this normatively loaded version of epistemic instrumentalism, and our solution to The Combinatorial Problem does not presuppose any particular version. But for ease of exposition, we shall work with an account that has been defended elsewhere by one of the authors.⁹

At the base of this account is a set of claims about how, exactly, our beliefs are normatively constrained by evidence. One suggestion would be to say that evidence for *p* always provides a *pro*

⁶ See e.g. Quine (1967).

⁷ For an influential version of this objection, see Kelly (2003); for a recent discussion, see Côté-Bouchard (2015).

⁸ E.g. Grimm (2009), Cowie (2014), Steglich-Petersen (2018).

⁹ Steglich-Petersen (2018).

tanto reason to believe p , and that decisive evidence for p makes it the case that one *ought* to believe p . But there are grounds for doubting that the normative significance of evidence is as strong as that. For one thing, if we consider all of the deductive consequences of what we know, we have decisive evidence for many more propositions than we could ever form beliefs about. Insofar as ‘ought’ implies ‘can,’ this would be a counterexample to the thesis that we ought to believe whatever we have decisive evidence for.¹⁰ But even if we *could* form all of the relevant beliefs, it is doubtful that we ought to do so. What is the point of believing a bunch of uninteresting tautologies? It seems more plausible to suppose that adequate evidence is required for the *permissibility* of belief, and that evidence only generates positive reasons to form beliefs about those propositions that there is practical value in having beliefs about.

If these normative relationships between evidence and belief are the proper explananda for theories of epistemic normativity, we can draw on general principles of instrumental normativity to explain them. In the literature on instrumental normativity, it is common to distinguish between two kinds of normative principles. The one kind concerns what is sometimes called ‘instrumental rationality,’ which is a matter of being *coherent* in one’s instrumental attitudes and actions. If one has a certain aim and believes some means to be necessary to achieve that aim, one is instrumentally irrational if one fails to either take the means or give up the aim. But this coherence norm of instrumental rationality is silent on what one has reason to do. Having instrumental reason to take the means to some aim requires that one has reason to adopt the aim in the first place. On one popular line of thought, instrumental reasons for taking means are generated through a process of ‘transmission’ of reasons from aims to means. In addition to the coherence norm of instrumental rationality, we thus need a principle for when and how instrumental reasons are transmitted in this way.¹¹

Together, these two kinds of principles provide the resources to explain why beliefs are normatively constrained by evidence the way they are. As shown elsewhere, the sense in which there generally is something faulty about evidentially unsupported beliefs can be explained in terms of the coherence norm of instrumental rationality.¹² And the sense in which evidence sometimes

¹⁰ For this point, see Whiting (2013).

¹¹ The exact shape and content of these principles is controversial. So is their compatibility. For the latter, see Setiya (2007).

¹² See Steglich-Petersen (2018, §4.1). In short, if we suppose that every belief is associated with an aim to believe the truth, and that believing in accordance with evidence is thought to be necessary to achieve this aim, it will be instrumentally incoherent not to intend to believe in accordance with the evidence. For defences of the required teleological account of belief, see Steglich-Petersen (2006; 2009).

provides positive reason to adopt a belief can be explained in terms of the principle of instrumental transmission. For our purposes, only the transmission principle will be relevant. This principle should capture the idea that one's reasons to take some means, *M*, depends on one's reasons for pursuing the aim, *A*, to which *M* contributes. And the *strength* of one's reasons to *M* should depend on the strength of one's reasons to pursue *A* as well as on the degree to which *M* contributes to achieving *A*. One way of capturing these dependencies, which avoids a number of standard problems and counterexamples, has been proposed by Niko Kolodny (2018):¹³

General Instrumental Transmission: If there is reason for one to pursue an aim, *A*, and there is positive probability conditional on one's *M*-ing, that this helps bring about *A* non-superfluously, then that is a reason for one to *M*, whose strength depends on the reason for one to pursue *A* and the probability.

When applied to *epistemic* reasons, the idea is that evidence for *p* speaks in favor of believing *p* only in contexts where there is a practical reason to pursue the aim of coming to a true belief as to whether *p*. With such a truth-aim in place, the evidence helps generate a reason to believe *p* by making it likely that doing so achieves the aim of coming to a true belief about *p*. If we make the suitable substitutions, we thus get the following instantiation of the general transmission principle:

Instrumental Transmission of Epistemic Reasons for Belief: If there is reason for one to pursue the aim of coming to a true belief about *p*, and there is positive probability conditional on one's adopting a belief that *p*, that this helps bring about that aim non-superfluously, then that is a reason for one to adopt a belief that *p*, whose strength depends on the reason for one to pursue the aim and the probability.¹⁴

This principle encodes the idea that when evidence for *p* provides positive normative reason to believe *p*, it does so because it raises the probability that believing *p* will result in a true belief about *p*, and because there is reason to pursue a true belief about *p*. When one has no reason to pursue this truth-aim, no amount of evidence for *p* will give rise to a normative reason to believe *p*.

What, on this view, should we make of the common practice of referring to evidence *simpliciter* as 'reasons' for belief? Since, on our view, evidence does not by itself constitute a

¹³ The terminology has been adjusted for our purposes.

¹⁴ We take the phrase 'the aim of coming to a true belief as to whether *p*' to refer to the dual aim of believing *p* when *p* is true, and not believing *p* when *p* is false. Later on, we will briefly discuss the relative weights of these aims, but for the most part, treating the dual aim as a single 'truth-aim' will do us no harm.

normative reason for belief, it is strictly speaking false to refer to one's evidence as 'reasons.' Rather, we think that this practice is best understood as a matter of talking about epistemic reasons in an 'elliptical' manner: we take the truth-aim for granted, and simply focus our attention on the question of which belief is most likely to achieve that aim in light of the available evidence. This phenomenon is also known from practical contexts in which some aim 'goes without saying.' Take the case of economic discourse: it is natural to speak of someone as having 'economic reason' to take some action that is likely to further some economic aim (e.g., the aim of becoming debt-free). But the mere fact that the action is likely to yield a given economic result does not *on its own* constitute a normative reason in its favor; it only does so in combination with a reason to pursue the economic aim to begin with.

The same goes for epistemic reasons. In fact, when it comes to epistemic reasons, we should expect elliptical statements to be even more common. This is because, with respect to the vast majority of propositions—and certainly the propositions that find their way into ordinary reasoning and conversation about what to believe—we have at least *some* practical reason to pursue the aim of believing what is true and to avoid believing what is false. Nevertheless, it can sometimes become relevant to draw attention to the question of whether there is reason to pursue the underlying truth-aim. To take a somewhat contrived example, imagine someone pointing out various trivialities that follow from things that you know, and that you have not yet formed any beliefs about. In such a case, you'd be right to object 'why should I care?' and proceed to ignore what you are told. It is natural to interpret this as a matter of rejecting that you have any reason to seek a true belief about the propositions presented to you. It is like rejecting the normative relevance of instrumental advice on the grounds that you have no reason to pursue the relevant end. But in most conversational contexts, it will be entirely felicitous, even on our view, to presuppose reasons for the relevant aim and speak of evidence *simpliciter* as reasons.

The instrumentalist account outlined here is compatible with a variety of views about when and why we have reason to pursue the aim of forming true beliefs and avoiding false ones. Some philosophers maintain that there is always some intrinsic value in believing what is true, and some intrinsic disvalue in believing what is false. While we are skeptical of this idea, it is in principle compatible with our instrumentalist picture, as long as these intrinsic values can be weighed against other non-alethic values. On our favored approach, however, we have reason to adopt the aim of forming true beliefs only insofar as doing so is likely to help achieve the practical aims we have reason to pursue. These practical aims need not be restricted to prudential interests, but might

concern moral, intellectual, and aesthetic matters as well. For example, even if it is not in my own narrow prudential interest to seek true beliefs about the environmental impact of my travel habits, there might be moral reasons to do so.

Our version of epistemic instrumentalism is also compatible with various more general normative frameworks. In particular, it is neutral between consequentialism and other rival theories, such as deontological ones. This is because it relies on general principles of instrumental rationality and reasons-transmission, which must presumably be part of *any* normative framework. So, while it is compatible with our view approach that reasons for pursuing true beliefs are based on consequentialist considerations about the promotion of value, they need not be.

How does the instrumentalist account relate to the debate between evidentialists and pragmatists? It does not fall squarely within either camp. Consider first the pragmatist view on which purely practical considerations can generate genuine reasons for belief. Although our version of epistemic instrumentalism does not imply that there are such practical reasons for belief, it is compatible with their existence. We can see this by reconsidering the transmission principle: imagine a situation where one has strong practical reason to pursue some valuable outcome (say, a large monetary reward), and the best means to that outcome is to believe that *p*, *regardless of its truth*. In such a case, the transmission principle implies that one has instrumental reason to form such a belief. This prompts the question: is the transmitted reason a reason for believing *per se*, as the pragmatists claim, or merely a reason to ‘bring about’ the belief? Epistemic instrumentalism is in principle silent on that question. We have sympathies for the pragmatist answer, but the matter is not settled by the instrumentalist account outlined here.

Consider next the evidentialist view on which all reasons for belief are epistemic in nature. On one interpretation, this view is incompatible with epistemic instrumentalism: if all and only considerations that bear directly on the truth of *p* can ever play a role in constituting a reason to believe *p*, our instrumentalist account fails in virtue of holding that evidence can never by itself constitute a complete reason to believe *p*. But there are other interpretations of evidentialism that still retain its spirit. Consider, for example, the view that having adequate evidence for *p* is merely *necessary* for having reason to believe *p*. Such a position is compatible with epistemic instrumentalism, as long as one never has genuine practical reasons for belief, which, as we have seen, is an issue that instrumentalism in principle stays silent on.

Before we shift our focus back on The Combinatorial Problem, it will be helpful to see how the instrumentalist account can explain the previously mentioned differences in the combinatorial

properties of practical and epistemic reasons. Consider first the *scales* on which the weights of epistemic and practical reasons appear. Are the weights of epistemic instrumental reasons bounded or unbounded? The answer depends on which aspect of these reasons we focus on. The reasons one might have for adopting the aim of believing the truth as to whether p are unbounded in weight: there is no apparent upper limit beyond which one cannot acquire any further reason to adopt the truth-aim. This is particularly clear if we think of these reasons as deriving from the practical importance of having true beliefs, in which case the difference between true and false beliefs might be the difference between life and death. By contrast, the probabilistic aspect of epistemic instrumental reasons is clearly bounded: no amount of evidence can make it more than 100% likely that some belief is true. Now, in most conversational contexts, we take the truth-aim for granted and focus solely on the bounded aspect of epistemic instrumental reasons, which we will talk about in an elliptical manner as if it were the whole reason. For example, when we say that there is ‘conclusive reason’ to believe that p , the non-elliptical rendering would be that there is ‘conclusive *evidence* for p , and hence strong reason to believe that p , given that there is reason to pursue the truth-aim.’ On the present picture, this is why epistemic reasons *seem* to have bounded weights, although, strictly speaking, they do not.

Much the same can be said about the combinatorial difference between prohibitive and permissive balancing. Again, we must pay attention to which aspect of instrumental epistemic reasons we focus on. The reasons one might have for adopting the aim of believing the truth about p exhibit permissive balancing, since, when one has equally strong reason for and against adopting the truth-aim, either option is permissible. By contrast, the probabilistic aspect of instrumental epistemic reasons exhibits prohibitive balancing. To see this, we will need to make a working assumption about the relative importance of *believing what is true* and *not believing what is false*. A natural suggestion might be to treat these aims as equally important. This would mean, on our view, that one should believe p whenever the evidence makes p more than 50% likely to be true, and that one should avoid believing p whenever the evidence makes p less than 50% likely to be true. But it would not explain why one should suspend judgment about p when the evidence makes p exactly 50% likely to be true, nor, indeed, why one should suspend judgment when the evidence only points weakly in either direction. To explain this, we must regard it as generally more important to avoid false beliefs than to secure true ones.¹⁵ Obviously, this is a substantive assumption that calls out for an explanation. But as we argue elsewhere, this is an explanatory

¹⁵ For an early discussion of this point, see James (1979). For a more recent discussion, see Horowitz (2018).

challenge that is shared among a wide range of epistemological theories.¹⁶ Assuming that the challenge can be met, the probabilistic aspect of instrumental epistemic reasons exhibit prohibitive balancing. And given that we tend to think and talk about such reasons in an elliptical manner, this type of balancing will normally be what stands out when we consider what there is reason to believe.

3. Epistemic Instrumentalism and The Combinatorial Problem

With this outline of epistemic instrumentalism in hand, we now turn to consider how it can solve The Combinatorial Problem. We begin by developing a diagnosis of the problem; a diagnosis, that is, of why practical and epistemic reasons seem, in Berker's words, like 'oil and water.' Our suggestion, in short, is that the apparent incomparability of practical and epistemic reasons stems from a subtle, but crucial, mistake of comparing *incomplete* instrumental reasons for belief. This diagnosis will then point towards a solution to The Combinatorial Problem. All of this will be unpacked in due course. But to prepare the ground, it will be useful to consider how instrumental reasons for *action* interact. Once we have the purely practical case in plain view, it will be easier to see how instrumental reasons for belief interact.

There are many different types of situations that call out for a comparison of opposing instrumental reasons for action. Most obviously, this happens when one has reason to pursue two aims that cannot both be realised, perhaps because the only two means to those aims cannot both be taken. For example, working overtime might be the only means to earning extra money, while *not* working overtime might be the only means to spending more time with the family. Here a weighing challenge arises, because the two means are mutually exclusive. But the need to compare instrumental reasons for action also arises in cases where the relevant aims merely *detract* from one another, or when the best means to the one aim somehow *limits* one's ability to take the means to the other aim. For example, the aim of increasing one's fitness level might detract from the aim of getting a promotion, because one only has limited time available to pursue those aims.

How might one go about comparing one's instrumental reasons for action in such cases? According to the transmission principle, the weight of any given instrumental reason is determined by two factors: (i) the weight of the reasons one has to pursue the relevant aim, and (ii) the likelihood that the relevant means will be successful in achieving that aim. Both determinants

¹⁶ Skipper & Steglich-Petersen (ms).

should thus be taken into account when assessing which action there is most instrumental reason to perform. Suppose I have much stronger reason to pursue Aim-1 than Aim-2, but that the best means to Aim-1 is very unlikely to achieve the aim, whereas the best means to Aim-2 will almost certainly achieve that aim. In at least some such cases, one should take the means to the less desirable aim on account of its more probable success. For example, I might prefer becoming the first 50-year-old world heavyweight champion over retaining my current job as a clerk, yet have very little instrumental reason to skive off work to spend six hours a day in the gym, because doing so would most likely not bring me the championship belt. The more general lesson here is that when comparing different instrumental reasons for action, it is not enough to focus on just *part* of those reasons; one must look at the entire set of factors that determine the weights of those reasons. That is what we mean by saying that one cannot meaningfully compare ‘incomplete’ instrumental reasons.

Mistakes of comparing incomplete instrumental reasons for action might manifest themselves in a number of different ways. One type of mistake is to compare the likelihood that Means-1 will achieve Aim-1 to the likelihood that Means-2 will achieve Aim-2, while ignoring the question of which aim there is most reason to pursue. Suppose I base my decision about whether to work overtime today solely on a comparison of the likelihood that doing so will get me a pay raise (say, 50%), and the likelihood that *not* doing so will leave more time to spend with my family (say, 99%). Such a comparison is, of course, possible to make. But to determine whether there is most instrumental reason to work overtime or not, I must also look at my reasons for pursuing each aim.

Another type of mistake is to compare one’s reasons to pursue Aim-1 with one’s reasons to pursue Aim-2, while ignoring the question of which means is more likely to achieve those aims. Suppose I base my decision about whether to stay at work or go to the gym solely on a comparison of my reasons for pursuing a promotion to my reasons for pursuing an increased fitness level. Again, such a comparison is fully possible to make. But to determine whether I have most instrumental reason to stay at work or go to the gym, I must also look at the likelihoods that those actions will be successful in achieving the respective aims.

Finally, a more subtle mistake is to compare the likelihood that Means-1 will achieve Aim-1 to the entire instrumental reason constituted by the likelihood that Means-2 will achieve Aim-2, and one’s reasons for pursuing Aim-2. Suppose I have strong instrumental reason to train for an ironman (say, because I want to be perceived as a successful athlete, and expect myself to be able to complete an ironman if I train for it), yet decide against it on the grounds that the life expectancy

of ironmen is low. This comparison does not settle the question of whether I have most instrumental reason to train for the ironman or not, because the answer to this question partly depends on how strong reasons I have to pursue a long life. I might prefer to live fast and die young, or to live slow and be able to see my grandchildren grow up.

The reason why each of these comparisons fail to deliver a definite verdict as to which action there is most instrumental reason to perform is that they ignore at least one of the determinants featuring in the transmission principle. Only when a *complete* instrumental reason for some action is compared to a *complete* instrumental reason for some other action do we get an answer to the question of what there is most instrumental reason to do.

The same goes for instrumental reasons for *belief*: only when a complete instrumental reason for some belief is compared to a complete instrumental reason for some other belief do we get a definite verdict as to which belief there is most instrumental reason to adopt. This is where our diagnosis of The Combinatorial Problem lies. When seen through the lens of epistemic instrumentalism, the apparent incomparability of epistemic and practical reasons stems from a mistake of treating evidence as if it on its own constitutes a complete reason for belief. For the most part, no harm is done by thinking and talking about epistemic reasons in this elliptical manner, and it is quite understandable that we are prone to forgetting the whole picture. But in cases where the truth-aim is called into question, we run into trouble if we continue to think about epistemic reasons in an elliptical manner. For in doing so, we easily fall prey to the mistake of comparing an incomplete instrumental reason for belief (such as Dr. Fisch's evidence about the winter weather) to a complete instrumental reason for belief (such as the reason constituted by Bill's offer and Dr. Fisch's reason to pursue a monetary gain). On our view, this is what explains the temptation to think that a solution to The Combinatorial Problem requires us to mix 'oil and water.'

That's the diagnosis. Now for the solution. On the present picture, how should we think about the interaction between Dr. Fisch's epistemic and practical reasons for belief? Let us begin by considering the epistemic reasons. We have assumed that Dr. Fisch's evidence makes it highly likely that the herring run will come early this year. On the instrumentalist account, this does not by itself provide a reason to believe that the herring run will come early; it only does so in combination with a reason to pursue the aim of coming to a true belief on the matter. Let us assume that Dr. Fisch has such a reason (say, based on a desire for a huge herring haul). Taken together, these two factors provide Dr. Fisch with an epistemic instrumental reason to form the belief that the herring run will come early, as specified by the transmission principle. Likewise, we

can understand Dr. Fisch's practical reasons for belief in terms of transmission. We have assumed that Dr. Fisch has reason to pursue the aim of winning the \$1M, and that Bill's offer makes it certain (or at least highly likely) that Dr. Fisch will receive the reward if he forms the belief that the herring run will come late. Taken together, these factors provide Dr. Fisch with a practical instrumental reason to form the belief that the herring run will come late.

Given that Dr. Fisch cannot both believe that the herring run will come early *and* late, he thus faces a task of weighing his instrumental reasons. But how? From the foregoing considerations it emerged that any comparison of instrumental reasons involves a comparison of one complete instrumental reason with another complete instrumental reason. Accordingly, Dr. Fisch must compare his complete epistemic instrumental reason to believe that the herring run will come early with his complete practical instrumental reason to believe that the herring run will come late. The weight of the epistemic instrumental reason is determined by the two factors in the transmission principle: Dr. Fisch's reason to pursue a true belief as to whether the herring run will come early or late, and the likelihood that he will end up with a true belief if he forms the belief that the herring run will come early. Likewise, the weight of the practical instrumental reason is determined by Dr. Fisch's reason to pursue the aim of winning \$1M, and the likelihood that he will receive the reward if he forms the belief that the herring run will come late. Once Dr. Fisch has determined the weight of each of these instrumental reasons, he should be in a position to determine which belief there is most instrumental reason to form.

So which belief *is* there most instrumental reason for Dr. Fisch to form? The answer depends on the details of the case. If Dr. Fisch is desperately in need of money, there might be most instrumental reason for him to form the belief that the herring run will come late. By contrast, if he is determined to win the title as 'fisherman of the year,' there might be most instrumental reason for him to believe that the herring run will come early. However we fill in the details, the instrumentalist account offers a way for Dr. Fisch to weigh his epistemic and practical reasons for belief in a way that closely resembles how we ordinarily weigh different instrumental reasons for action. This is not to say that it won't be challenging for Dr. Fisch to determine what he has most instrumental reason to believe. But the challenge is no harder than those arising whenever practical reasons for action must be weighed against each other.

The key to our solution of The Combinatorial Problem is that the process of weighing Dr. Fisch's epistemic and practical instrumental reasons for belief does not involve any direct comparison of his evidence to his practical reasons. The apparent need for such a comparison is

what made The Combinatorial Problem look intractable to begin with. But if we accept the instrumentalist account of reasons for belief, no such comparison is needed. The temptation to think otherwise stems from an ingrained habit of thinking and talking about epistemic reasons in an elliptical manner. Once we remind ourselves that evidence does not by itself constitute a complete reason for belief, but only does so in combination with a reason to pursue a true belief, it becomes clear that The Combinatorial Problem can be solved without ever mixing ‘oil and water.’

4. Are Epistemic Reasons Outweighed Too Easily?

An objection that might be raised against our account is that it makes it too easy for one’s practical reasons to outweigh one’s epistemic reasons. The force of this objection obviously depends on one’s view about when the practical benefits of a false belief may outweigh the benefits of believing the truth. Extant discussions of this issue focus largely on cases with very significant practical benefits, such as the prospect of eternal bliss (Pascal 1670). Most would presumably agree that if practical reasons have any grip of our beliefs at all, extreme benefits such as these are likely to outweigh one’s epistemic reasons.

But what if the practical benefits are more modest? For example, can just a few dollars ever outweigh one’s epistemic reasons to believe the truth? On our account, this is possible, at least in principle. If the proposition in question is sufficiently uninteresting, even a very small monetary gain may be enough to tip the scale in favor of an evidentially unsupported belief. This result might violate the intuitions of some philosophers, even those of a pragmatic bent. Reisner (2008), for one, designs his ‘defeating function’ so that practical reasons can trump epistemic reasons only when the practical stakes are quite significant. Whenever the practical reasons fall below some relatively high threshold, the epistemic reasons win the day.

However, although our account makes it in principle possible for a practical reason of modest strength to outweigh one’s epistemic reasons, this will rarely be the case in practice. There are two reasons for this. First, even when the practical value of having some false belief would be greater than that of a true one, the costs associated with bringing about the false belief will often tip the balance in favor of pursuing the true belief. It is typically not a trivial matter to bring oneself to believe something against the evidence. After determining that believing in God would be in one’s best interest, Pascal (1670) famously went on to recommend various means to bring about such a belief, such as attending to mass regularly, which is clearly something that must be factored into

one's decision. If Pascal is right, the potential infinite benefit of believing in God greatly outweighs both the benefits of believing otherwise, and the cost of going to mass. But in more mundane cases, it seems that the psychological toil of bringing oneself to believe against the evidence could easily outweigh the practical benefits.

Secondly, since many of our beliefs are closely interwoven, a false belief can easily lead to other false beliefs, thus creating a risk of frustrating our aims, whether practical or epistemic, even if the initial false belief taken in isolation does not.¹⁷ And the *process* of causing oneself to adopt some belief against the evidence can in itself have harmful effects. From nature, we are equipped with dispositions and mechanisms that fairly reliably result in true beliefs.¹⁸ These are reinforced by acquired reasoning skills and epistemic habits. For the most part, these work in our favor, even if we are occasionally better served by false or evidentially ungrounded beliefs. But if we must tamper with these dispositions to bring about false or evidentially unsupported beliefs, there is a risk of making the dispositions overall less reliable. This is a further potential cost that must be taken into account when considering whether to follow one's practical reasons for belief.

Yet, our account remains committed to the claim that in those rare scenarios where one can easily bring oneself to adopt a belief against one's evidence, and very little harm would result from believing something false, even modest practical benefits may tip the balance in favor of forming that belief. This is a result that we are prepared to embrace. It does not seem implausible to us that practical considerations may outweigh the aim of believing the truth under such circumstances.

5. Is 'Double-Weighing' Troublesome?

Recently, Berker (2018) has raised a general objection against a class of views—what he calls 'double-weighing' views—to which our instrumentalist account can be seen as belonging. The idea behind the double-weighing view, as Berker construes it, is to apply a two-step procedure to determine which doxastic attitude towards some proposition is best supported by a given set of epistemic and practical reasons. The first step consists in weighing the epistemic reasons against each other in a way that exhibits prohibitive balancing. This results in a preliminary verdict about which doxastic attitude there is most epistemic reason to adopt. This attitude is then associated

¹⁷ For a related point, see Schroeder (2007).

¹⁸ Is there a tension between beliefs being 'from nature' sensitive to the truth, and our characterisation of the truth-aim as something that agents can choose not to pursue? We think not. It is commonplace that we can choose to act in ways that counteract certain natural dispositions. For example, we are naturally disposed to shy away from fire, but can choose to walk on it if we see reason to do so. See Steglich-Petersen (2009) for discussion.

with a normative reason in its favor, whose strength depends on the overall balance of the epistemic reasons. In the second step, this normative reason is weighed against the practical reasons in a way that exhibits permissive balancing, thus resulting in an overall verdict about which doxastic attitude there is most reason to adopt. The supposed benefit of this two-step procedure is that it allows us to combine epistemic and practical reasons while respecting their different combinatorial properties.

Our instrumentalist account can be construed as a version of the double-weighting view, where the first step consists in determining the weight of one's epistemic instrumental reasons for and against different candidate attitudes, and the second step consists in weighing these epistemic instrumental reasons against one's practical instrumental reasons for and against those attitudes. Of course, this construal may seem somewhat artificial, since there is not inherent order of the steps on our view. But the construal makes it easier to see the connection to Berker's discussion.

According to Berker, any double-weighting view faces a challenge brought out by comparing the following cases:

- (a) Something terrible will happen if I believe p , and I have some evidence against p and some equally strong evidence for p (and no other evidence bearing on p).
- (b) Something terrible will happen if I believe p , and I have some evidence against p and some much stronger evidence for p (and no other evidence bearing on p).

The only difference between these cases is that my evidence for p is stronger in the latter than in the former. Thus, if disbelieving p is forbidden in case (a), disbelieving p must, *a fortiori*, be forbidden in case (b) as well. However, according to Berker, the double-weighting view does not deliver this result. In case (a), the first step results in a normative reason to suspend judgment about p , since the evidence combines in a prohibitive manner. This reason is then, in the second step, combined with a very strong practical reason to either disbelieve or suspend judgment about p . Thus, I have most reason to suspend judgment about p , which means that I'm forbidden to disbelieve p . In case (b), the first step instead results in a normative reason to believe p . This reason is then, in the second step, combined with a very strong practical reason to either disbelieve or suspend judgment about p . Thus, I have just as much reason to disbelieve p as to suspend judgment about p , which means that I'm permitted to disbelieve p .

According to Berker, these are 'excessively odd verdicts:' surely *gaining* evidence for p should not have the effect of *elevating* the normative status of disbelieving p (2018, p. 455). Here we agree.

We also agree that the double-weighting view, as Berker construes it, delivers these implausible verdicts. But we do not think that Berker's construal of the double-weighting view is the most charitable one. Specifically, our instrumentalist account provides a better construal of the *first* step. As Berker describes the first step, its aim is to determine a 'winner,' as he calls it, among the candidate doxastic attitudes, from an epistemic perspective. The winner is then associated with a reason in its favor, whereas the losers are not. No consolation prizes are awarded. Only the winner enters the second step with any normative support from the evidence.

The lack of consolation prizes is effectively what generates the implausible verdicts. Consider what happens if we suppose that every candidate doxastic attitude goes away from the first step with a reason in its favor, whose strength depends on how well the attitude is supported by the evidence. This gives us a way of differentiating among losers, intuitively speaking. For example, when the evidence speaks in favor of believing p , it seems reasonable to suppose that disbelieving p enjoys less evidential support than does suspending judgment about p . By contrast, when the evidence is perfectly balanced for and against p , it seems that believing p and disbelieving p is equally unsupported by the evidence.

Once the double-weighting view is construed so as to award consolation prizes, we get the right verdicts in Berker's cases. In case (a), the first step still results in a reason to suspend judgment about p , but now *also* results in two equally much weaker reasons to believe and disbelieve p respectively. These reasons are then, in the second step, combined with the very strong practical reason to either disbelieve or suspend judgment about p . Thus, I have most reason to suspend judgment about p , which means that I'm forbidden to disbelieve p . In case (b), the first step still results in a reason to believe p , but now *also* results in two weaker reasons to disbelieve and suspend judgment about p respectively, where the latter reason is stronger than the former. These reasons are then, in the second step, combined with the very strong practical reason to either disbelieve or suspend judgment about p . Thus, I have most reason to suspend judgment about p , which means that I'm forbidden to disbelieve p . So, on this construal of the double-weighting view, I'm forbidden to disbelieve p in *both* of Berker's cases. These verdicts, we take it, do not seem implausible or odd.

Our instrumentalist account can be seen as a consolation-prize-awarding double-weighting view. The mere fact that some means to an aim is more likely to be successful than any other means does not entail that there is *no* instrumental reason to take the other possible means. All it entails is that there is *less* instrumental reason to take those means. In particular, when it comes to

instrumental reasons for *belief*, the fact that believing p is the attitude most likely to achieve the aim of coming to a true belief about p does not entail that there is *no* instrumental reason to take any other candidate attitudes, or that these other attitudes must be *equally* unsupported. This is what allows us to steer clear of Berker's objection. Consider how our view applies to the cases from above. In case (a), the first step delivers the verdict that there is most epistemic instrumental reason to suspend judgment about p , and that there is equally little epistemic instrumental reason to believe p as to disbelieve p . These epistemic instrumental reasons are then, in the second step, combined with the very strong practical instrumental reason to either disbelieve or suspend judgment about p . Thus, I have most reason to suspend judgment about p , which means that I'm forbidden to disbelieve p . In case (b), the first step delivers the verdict that there is most epistemic instrumental reason to believe p , and that there is more epistemic instrumental reason to suspend judgment about p than to disbelieve p . These epistemic instrumental reasons are again, in the second step, combined with the very strong practical instrumental reason to either disbelieve or suspend judgment about p . Thus, I have most reason to suspend judgment about p , which means that I'm forbidden to disbelieve p . In sum, our instrumentalist account renders disbelief in p forbidden in both of Berker's cases. This shows that our account does not suffer from the extensional inadequacies that Berker raises against double-weighting views.

6. Conclusion

We have argued that epistemic instrumentalists are in a good position to solve what we called 'The Combinatorial Problem': the problem of how to weigh epistemic and (putative) practical reasons for belief. By understanding epistemic reasons as a species of instrumental reasons, we argued that The Combinatorial Problem can be reduced to a relatively benign problem of how to weigh different instrumental reasons. In addition, the instrumentalist account has a number of further theoretical and explanatory attractions. First, the account derives from an independently motivated transmission-based account of how different instrumental reasons for belief are to be weighed and compared. This sets the account apart from gerrymandered weighing functions that are reverse-engineered to encode certain intuitive verdicts about how epistemic and practical reasons for belief interact. Second, the instrumentalist account has unificatory power in virtue of being embedded in a single, unified theory of epistemic and practical normativity. Finally, the instrumentalist account has a kind of 'diagnostic' value in virtue of providing an explanation of the

apparent intractability of The Combinatorial Problem in terms of a common tendency to think and talk about epistemic reasons in an ‘elliptical’ manner. Thus, epistemic instrumentalism not only provides the resources to solve The Combinatorial Problem without ‘mixing oil and water.’ It also explains the temptation to think that such mixing is called for in the first place.

References

- Adler, J. (2002). *Belief's Own Ethics*. MIT Press.
- Berker, S. (2013). ‘Epistemic Teleology and the Separateness of Propositions.’ *Philosophical Review*, 122: 337-393.
- (2018). ‘A Combinatorial Argument Against Practical Reasons for Belief.’ *Analytic Philosophy*, 59: 427-470.
- Booth, T. (2012). ‘All Things Considered Duties to Believe.’ *Synthese*, 187: 509-517.
- Côté-Bouchard, C. (2015). ‘Epistemic Instrumentalism and the Too Few Reasons Objection.’ In: *International Journal of Philosophical Studies*, 23: 337-355.
- Cowie, C. (2014). ‘In Defence of Instrumentalism about Epistemic Normativity.’ *Philosophical Studies*, 191: 4003-4017.
- Feldman, R. (2000). ‘The Ethics of Belief.’ *Philosophy and Phenomenological Research*, 60: 667–695.
- Grimm, S. (2009). ‘Epistemic Normativity.’ In: *Epistemic Value*, A. Haddock, A. Millar and D. Pritchard (eds.), Oxford University Press.
- Horowitz, S. (2018). ‘Epistemic Value and the Jamesian Goal.’ In: *Epistemic Consequentialism*, K. Ahlstrom-Vij and J. Dunn (eds.), Oxford University Press.
- James, W. (1979). *The Will to Believe and Other Essays in Popular Philosophy*. Harvard University Press.
- Kavka, G. (1983). ‘The Toxin Puzzle.’ *Analysis*, 43: 33-36.
- Kelly, T. (2003). ‘Epistemic Rationality as Instrumental Rationality: A Critique.’ *Philosophy and Phenomenological Research*, 66: 612-40.
- Kolodny, N. (2018). ‘Instrumental Reasons.’ In: *The Oxford Handbook of Reasons and Normativity*, D. Star (ed.), Oxford University Press.
- McCormick, M. (2015). *Believing Against the Evidence: Agency and the Ethics of Belief*. Routledge.
- Pascal, B. (1670 [1995]). *Pensées*. A. Kreilsheimer (ed.), Penguin Press.

- Quine, W. V. (1967). 'Epistemology Naturalized.' In: *Ontological Relativity and Other Essays*. Columbia University Press.
- Reisner, A. (2008). 'Weighing Pragmatic and Evidential Reasons for Belief.' *Philosophical Studies*, 138: 17-27.
- (2009). 'The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem.' *Philosophical Studies*, 145: 257-72.
- (ms.). 'Comparing Pragmatic and Alethic Reasons' Chapter II of *The Pragmatic Foundations of Theoretical Reason* (unpublished manuscript).
- Rinard, S. (forthcoming). 'Believing for Practical Reasons.' *Nous*.
- Schroeder, M. (2007). *Slaves of the Passions*. Oxford University Press.
- Skipper, M. & A. Steglich-Petersen (ms). 'On Belief Gambles.' Unpublished manuscript.
- Steglich-Petersen, A. (2018): 'Epistemic Instrumentalism, Permission, and Reasons for Belief. In: *Normativity: Epistemic and Practical*, C. McHugh, J. Way & D. Whiting (eds.), Oxford University Press.
- (2009). 'Weighing the Aim of Belief.' *Philosophical Studies*, 145: 395-405.
- (2006). "No Norm Needed: On the Aim of Belief." *The Philosophical Quarterly*, 56: 499-516.
- Setiya, K. (2007). 'Cognitivism About Instrumental Reason.' *Ethics*, 117: 649-673.
- Shah, N. (2006). 'A New Argument for Evidentialism.' *The Philosophical Quarterly*, 56: 481-498.
- Whiting, D. (2013). 'Truth: The Aim and Norm of Belief.' *Teorema*, 32: 121-136.