

ARTICLE

On the Practical Significance of Irrelevant Factors

Seyed Mohammad Yarandi

University of Toronto Scarborough Department of Philosophy Toronto, Ontario, Canada
Email: seyed.yarandi@mail.utoronto.ca

Abstract

I focus on an overlooked aspect of the challenge of irrelevant influences. The challenge is often framed in terms of whether recognizing the presence of irrelevant factors in the pedigree of a belief provides a defeater. I argue that the epistemic significance of irrelevant factors goes beyond their status as defeaters. I focus on what I call gray cases, where learning about such factors causes epistemic worry without justifying giving up the belief. I argue that in gray cases, the subject finds a commendatory epistemic practical reason to engage in an activity intended to assuage that epistemic worry.

Keywords: Irrelevant influences; epistemic reasons for action; higher-order evidence; belief revision; commendatory reasons; nonrequiring reasons

A Persian friend once told me that the students of sociology at her university are deeply skeptical of the International Monetary Fund's recommendations for developing countries. By contrast, the students of economics take the recommendations to be sound and defensible. The two groups often fight tooth and nail to convince each other, though their efforts seldom meet with success. I remarked that perhaps the students of economics are better positioned to make such an assessment. She denied that there is any difference in terms of expertise that helps one group to understand or assess the recommendations better. The problem, she maintained, is that both groups are indoctrinated by their professors, and then added, perhaps jokingly, that she would take the judgment of a philosopher more seriously than a member of either group.

Unfortunately, I had no specific opinion on the International Monetary Fund's recommendations. But her story exemplifies a type of concern in epistemology about which I do have an opinion, namely, the 'challenge of irrelevant influences.' The challenge can be formulated as follows: How should a subject respond upon identifying an 'irrelevant factor' in the pedigree of her beliefs? An irrelevant factor is a one that plays a constitutive role in the process of belief formation without being connected to its veracity.

It is common to challenge a subject for holding a belief by pointing out a consideration that shows the influence of an irrelevant factor (e.g., the subject's upbringing, his religious/political affiliation, or even his social class) in the formation of his belief. These factors show up in the causal explanation of why the subject holds a belief but does not justify it. For instance, consider criticisms like "Max finds Dummett's view plausible because he is an Oxford graduate!" or "Sarah buys Marx's argument in the first part of Grundrisse since she is a member of the communist party." These criticisms put forward a consideration that alludes to the presence of an irrelevant factor in the pedigree of the targeted belief. The presence of such a factor in a belief's pedigree sometimes comprises not just how the belief was formed, but also how it was maintained.

Intuitively, this type of consideration (henceforth, ‘IF considerations’) can instigate epistemic worry in us. But it is not clear how we should respond to IF considerations.

What distinguishes an IF consideration from familiar skeptical concerns is that it does not just point to the possibility of error (e.g., it is possible that the subject is just a brain in the vat); an IF consideration implies that error is not just a possibility but a genuine probability. Despite this difference, some philosophers have argued that in terms of the consequences, accepting the epistemic significance of IF considerations can be as detrimental as accepting skeptical scenarios (White 2010): that is, we should give up many of our intuitively innocuous beliefs. Hence much may depend on the outcome of this debate.

Though philosophers disagree on the epistemic significance of IF considerations, there is a consensus on how to understand them: they supposedly provide a defeating piece of higher-order evidence that counts against the credibility of the agent’s belief. I call this view the ‘doxastic interpretation of IF considerations.’ The disagreement is about whether such a consideration ever performs its purported function—i.e., whether an IF consideration ever provides a reason to give up the impacted belief. Philosophers are divided on this question. The naysayers deny that IF considerations provide a defeater and dismiss the challenge of irrelevant influences as just another version of skeptical challenges.¹ The yeasayers claim that at least in some cases, becoming aware of an irrelevant factor provides the subject with a defeater.²

In this paper, I will argue that to resolve the disagreement, we first need to reconsider the point of consensus—i.e., the doxastic interpretation of IF considerations. My claim is that a proper response to an epistemically significant IF consideration does not always involve revising the impacted belief. Sometimes, such a consideration can have a practical significance: it provides a practical reason to engage in an epistemic activity. I will argue that in such cases, the IF consideration does not provide us with a sufficient reason to revise the impacted belief. Still, there are cases in which IF considerations rationally compel us to make such a revision. The lesson to draw from these observations is simple: not all epistemically significant IF considerations demand the same kind of response. Therefore, we should abandon the doxastic interpretation for a position that I call the ‘hybrid interpretation of IF considerations.’

The claim that IF considerations sometimes provide us with a practical reason invites a challenging question: Is such a reason also an epistemic reason? I will argue that the answer to this question must be positive, and, therefore, I will demonstrate that there are epistemic reasons for actions. As a result, my defense of the hybrid interpretation also involves arguing for a distinction between two kinds of epistemic reasons: ‘epistemic doxastic reasons’ and ‘epistemic practical reasons.’ Those epistemic reasons that count in favor of believing, disbelieving, and suspending judgment belong to the first group—i.e., epistemic doxastic reasons. But, as I will argue, there are also epistemic reasons that favor engaging in certain epistemic activities—e.g., rechecking the veracity of some beliefs or gathering more first-order evidence. I call those reasons epistemic practical reasons.

The very idea of epistemic practical reasons can raise some eyebrows. Does accepting epistemic practical reasons commit us to the much-criticized idea that subjects are sometimes epistemically required to engage in an epistemic activity? My answer to this question is firmly negative. The reason for this is that I take epistemic practical reasons to be commendatory: they commend engaging in a course of action without issuing any deontic directive. Therefore, they do not provide

¹The naysayers often try to explain away the epistemic significance of those cases of irrelevant influences intuitively involving a defeater by showing that they are also an instance of another familiar epistemological problem, such as the problem of peer disagreement. This position is forcefully defended in White (2010). His view was further defended in Mogensen (2015). For other defenses of the naysayers’ position, see Elga (2008) and Dworkin (1996, 2011).

²Yeasayers often claim that the subject needs to revise only in certain types of cases, say nonpermissive cases. But their positions differ on which types of cases require belief revision (see Avnur and Scott-Kakures 2015; Schoenfield 2014; and Vavova 2018).

basis for any epistemic duty and defying them is not grounds for any type of epistemic blame or criticism.³

Recognizing the possibility of epistemic practical reasons can help us better understand the epistemic significance of IF considerations. It also motivates a slight reformulation of the challenge of irrelevant influences. The current formulation reflects the presupposition that the epistemic significance of an IF consideration is limited to determining what we ought or ought not to believe. In some cases, however, an IF consideration can provide a reason that lacks any deontic force. Therefore, a better formulation of the challenge goes as follows: How *can* we respond if we identify an irrelevant factor in the pedigree of our beliefs? Note that in addition to its descriptive use, “can” also has a normative use as we see in a sentence like “Can I come in?” or “How can I help?” When we ask a normative question with “can,” the answer includes but is not limited to our obligations and duties. Additionally, the answer covers what we are permitted or commended to do. Therefore, it is better to formulate the central question of this paper with “can” rather than “should” or “ought.”

The outline of my answer to the challenge of irrelevant influences goes as follows. Identifying the presence of an irrelevant factor is epistemically significant if and only if it instigates an epistemic worry. In such cases, I argue that the relevant IF consideration can sometimes provide us with an epistemic doxastic reason to revise the impacted belief. But there are cases in which an IF consideration does not have a sufficient normative force to justify revising the impacted belief and yet it does instigate a degree of epistemic worry. As a result, the subject finds an epistemic practical reason to engage in epistemic activities aiming at assuaging an epistemic worry. After defending my answer, I will describe three types of activity that can ease a worry caused by the presence of an irrelevant influence.

The paper is structured as follows. In [section 1](#), I will distinguish three types of cases in which an irrelevant factor plays a role in the formation of a belief. In [section 2](#), I will focus on the notion of epistemic practical reasons. Our discussion in this section paves the way for [section 3](#), in which I will offer the hybrid interpretation of IF considerations. Finally, I address the main question of this paper in [section 4](#) by clarifying how a subject can respond if he identifies an irrelevant factor in the pedigree of her belief.

1. The demarcation worry and the challenge of irrelevant influences

In this section, my aim is to explain a reason why philosophers have found it difficult to reach a consensus on how we should respond to the challenge of irrelevant influences. A major source of the disagreement is that philosophers have failed to provide a principled way to distinguishing three types of scenarios. After setting out these three types and giving an example for each, I will explain why finding a principled way to distinguish them is important in addressing the challenge of irrelevant influences.

The function of my examples is to give a taste of what an instance of each type of scenario might be like. This will later help me to offer a principled distinction between the three types of scenarios. Consequently, I have tried to select an example for each type that intuitively belongs to it, yet I do not claim that no disagreement on my choice of examples is possible. After all, an example might intuitively belong to one of the categories, though we might later become convinced that it is actually an instance of another. Still, my examples serve their required function if they help us to see what an instance of each type might look like.

³An objector might point out that what I call doxastic epistemic reason does not require a subject to form any belief-attitude. In this sense, there is no difference between epistemic practical and epistemic doxastic reasons. However, epistemic doxastic reasons can license an epistemic criticism of the subject. This is not the case with epistemic practical reasons as I understand them.

An example of the first type of scenario involves the presence of an irrelevant factor whose discovery is sufficient to justify an immediate revision in the impacted belief. Call these scenarios ‘the doomed cases.’ For instance, assume that I discover that my belief in emotivism was formed under the influence of a highly skilled hypnotist who has strongly biased me toward Ayer’s arguments. At least *prima facie*,⁴ identifying this influence provides me with a defeater for my belief about emotivism.

In the second type of scenario, the presence of an irrelevant factor intuitively causes some degree of epistemic worry. Still, it is not clear what the right response to such criticism should be. These cases are perhaps the most challenging of the three and therefore, I will present a detailed example of this kind of case. I hope this difference in detail will not prove distracting. This is the example that we will return to frequently in the paper.

The Medical Historian: Silvia is a well-known historian of medicine. Though she is the daughter of a prominent member of the communist party, she describes herself as apolitical. She joins a research group that is investigating the contributions and successes of the Soviets in the field of medicine. Her group finds a set of confidential documents showing that in the late thirties, the Soviets faced a breakout of plague on their territory but managed to quarantine the infected individuals and so prevent an epidemic. Thrilled by her finding, she shares them with her friend Sarah. Without reading the documents, Sarah questions Silvia’s belief by pointing out that her upbringing in a communist household might have instilled in her an unconscious affinity for communists. Therefore, she should not trust her judgment about the documents.⁵

The irrelevant factor in the above case is Silvia’s upbringing in a communist household, which could play an epistemically worrying role in the formation of her belief. The presence of this factor is the basis for an IF consideration: “You only believe this because of your upbringing in a communist household.” However, despite being epistemically worrying, it is not clear that the IF consideration is sufficient to rationally compel Silvia to revise her belief. This might have been the case if she were a passionate communist herself. But we do not expect, for instance, an atheist researcher to give up her belief about the philosophical achievements of Avicenna just because he was born in a Muslim household. Yet it is hard to deny that Silvia’s upbringing might have negatively influenced her ability to assess the documents. This is epistemically worrisome. How should she then respond to Sarah’s IF consideration?

This case is representative of a class of examples that I call ‘gray cases.’ In a gray case, it is not clear whether the subject needs to revise her belief, though we find the presence of an irrelevant factor epistemically worrisome. We’ll be especially interested in gray cases.

Finally, there are scenarios in which an irrelevant factor plays a role in the formation of a belief, but we do not take its presence as instigating any sort of epistemic worry. For instance, to instill scientific curiosity in my child, I can tell him that he will be rewarded with an expensive toy if he learns how the litmus test works. After working with the litmus test for a time, he learns that vinegar turns the litmus test red. In this case, my child’s desire for an expensive toy plays a role in the formation of his belief that vinegar turns the litmus test red. Yet this role is innocuous and does not intuitively prompt any sort of epistemic worry. I call such scenarios ‘innocuous cases.’

Arguably, innocuous cases bear a structural similarity with the other two types of cases. But the proper response to these cases cannot possibly be the same. Thus, we need to account for the difference between these three types of cases. This is what I call the ‘demarcation worry in the case of irrelevant influences,’ or for short, the ‘demarcation worry.’

⁴In a situation in which I know that my belief is also based on some independent evidence, this might not provide any defeater.

⁵The story about the plague epidemic is loosely based on *Just the Plague* by Lyudmila Ulitskaya.

Part of the solution to the challenge of irrelevant influences lies in a satisfactory treatment of the demarcation worry. This becomes clear if we take a closer look at the arguments offered by the two sides of the debate. Philosophers often provide us with example-driven arguments showing that in many cases, we do find the presence of irrelevant factors to be deeply unsettling. Here, the favorite examples tend to come from doomed cases (see Schoenfield 2014, 203). These cases are allegedly sufficient to teach us a lesson: at least in some cases, identifying an irrelevant factor rationally requires us to revise the impacted belief. On the other hand, those who are skeptical about the epistemic significance of irrelevant factors point out that given the structural similarity between innocuous cases and the other two types of cases, we should embrace the same position regarding all cases. Thus, if we think the mere presence of an irrelevant factor gives us a reason for belief revision, we should accept that even the beliefs formed in an innocuous case are epistemically problematic and revise them. By going down this path, we have no choice but to give up a vast number of our beliefs. This is indeed an implausible result, and is reminiscent of radical skepticism (White 2010, 574–75; Mogensen 2015, 594–95).

Given our discussion, addressing the demarcation worry is a key to solving the challenge of irrelevant influences. The reason for this is simple: identifying the presence of an irrelevant factor does not call for the same response in all three types of cases. However, to explain the difference between these three types of cases, we first need to discuss the distinction between epistemic practical and epistemic doxastic reasons.

2. Epistemic reasons and practical recommendations

Our goal in this section is to develop a conceptual framework that enables us to address the demarcation worry. To this end, I argue for a distinction between ‘epistemic doxastic reasons’ and ‘epistemic practical reasons.’ I defend the controversial thesis that some practical reasons deserve to be a member of the very exclusivist club of epistemic reasons. However, such epistemic practical reasons cannot provide the grounds for any sort of epistemic duty. In this sense, they are similar to the commendatory moral reasons associated with supererogatory acts. They commend certain epistemic practices without issuing any deontic directive. Subsequently, I will argue that the same consideration that provides us with reasons to form or drop doxastic states can sometimes provide us with epistemic practical reasons. This observation, I will argue in the next section, is the key to addressing the demarcation worry and, eventually, the challenge of irrelevant influences.

Let us start with some preliminary facts about reasons. Reasons can be understood as considerations⁶ that count in favor of actions or beliefs.⁷ Thus, a given consideration becomes a reason in virtue of being in a ‘counting-in-favor relation’ with beliefs or actions.⁸ The same consideration can be in such a relation with different actions and beliefs. For instance, the consideration that I feel very hungry can count in favor of finishing a conversation sooner, rushing to a restaurant, or believing that I have not eaten much in the last hour or so. Therefore, a single consideration can provide us with many different reasons (Hieronymi 2005, 439).

Following Mark Schroeder, I understand epistemic reasons to be “whatever reasons bear on epistemic rationality” (2021, 145) ‘Epistemic doxastic reasons,’ often taken to be the only type of epistemic reasons, count in favor of one of the three paradigmatic doxastic states: belief, disbelief, and withholding belief. Thus, an essential feature of any epistemic doxastic reason is to support a doxastic recommendation to form one of the aforementioned states. By contrast, practical reasons (I mean practical reasons in general) count in favor of actions by rendering them rationally permissible, choiceworthy, or required.

⁶Regarding the ontology of reasons, I am more inclined to take them as true propositions (see Lord 2018).

⁷This analysis of reasons is put forward by Scanlon in his 1998 book.

⁸For a critical discussion of understanding reasons in terms of counting in favor relations, see Schroeder (2021, 42).

Practical reasons come in two types: ‘peremptory’ and ‘commendatory.’ The staple of a peremptory reason is that it comes with a deontic charge and therefore, in the absence of countervailing reasons, defying it provides ground for criticism. By contrast, a commendatory reason commends an action without issuing a deontic directive.⁹ Such a reason has a justificatory force but lacks a requiring power (Scanlon 2014, 107). Thus, in the absence of countervailing reasons, having a commendatory reason for ϕ ing renders ϕ ing *choiceworthy* without mandating it.

There is a well-respected orthodoxy among epistemologists that no one can be epistemically blameworthy or criticizable for not engaging in an epistemic practice such as evidence gathering. The reason is that a subject’s epistemic duties only concern how he is supposed to respond to the evidence he already possesses.¹⁰ Our epistemic duties are simply silent about whether one ought to engage in any course of action. From this observation, many hasten to conclude that there is no such thing as an *epistemic practical reason*—i.e., epistemic reasons that favor engaging in a course of action.¹¹ Others have gone to the other extreme by questioning the plausibility of the aforementioned orthodoxy.¹² Both extremes seem to ignore the possibility of some plausible wiggle room, I believe.

A way to reconcile orthodoxy with accepting epistemic practical reasons is to deny that such reasons can provide the basis for any epistemic duty. More precisely, we can take epistemic practical reasons to be exclusively commendatory—i.e., they provide normative support for a course of action without issuing a deontic directive to perform it. In this way, we can preserve the orthodoxy according to which a subject is not epistemically required to perform any action while endorsing the idea that some epistemic reasons can favor actions.

But the above solution at best shows that epistemic practical reasons are possible, not that they actually exist. To argue for their existence, I need to show that there are reasons that can be legitimately called epistemic and practical at the same time. Consider the following scenario:

Library: After two months of research on historical documents in a library, Jane comes to an interesting conclusion about the daily life of Catherine the Great: she used to sing Homer’s poems while taking baths. Jane shares this interesting data with her historian friends. After a while, Sam informs her that some of the historical documents in one of the libraries in the city were replaced by forged misleading documents. The city has more than eight hundred libraries and Sam does not know in which of the eight hundred the forgery happened. Subsequently, Jane calls the library, and the receptionist tells her that she cannot confirm whether the forgery includes her research material. However, she can confirm that they have decided to close the historical documents section since more than ninety percent of the historical documents in that library had been replaced by forged documents last year (i.e., six months before Jane started her research).

Let us focus on the first true proposition that Jane comes to know about the incident: a forgery has happened in one of the eight hundred libraries in the city. We can distinguish two questions related to the normative significance of this piece of information. First, can Jane cite the very fact that a forgery has happened in a library in the city to justify giving up her belief? Second, can Jane cite the same fact to justify her attempt to find out more about the incident?

⁹See Little and Macnamara (2022), Horgan and Timmons (2010), and Dancy (2004). For an overview of the literature on nonrequiring reasons, see Little and Macnamara (2020).; For the purpose of this paper, I use “nonrequiring reasons” in a narrow sense as referring to those reasons that commend an action without issuing a deontic directive. But there is another type of nonrequiring reason called a ‘permissibility-conferring reason’ that renders an action permissible without commending it (see Gert 2004). See Little and Macnamara (2020) for a helpful discussion of these two types of nonrequiring reasons.

¹⁰For a forceful defense of this view see Feldman (2000, 2002). See also Kelly (2003) and Nelson (2010).

¹¹See Kelly (2003) and Hedden (2015). Pamela Hieronymi criticizes the idea of epistemic practical reasons based on a different consideration (2005). I address her criticism in footnote 14.

¹²See Flores and Woodard (2023) and Peels (2016). For a different take on the issue see Booth (2006).

I believe the answer to the first question is negative. Though the very fact that a forgery has happened in one of the eight hundred libraries in the city puts some limited pressure for revision on Jane's belief, the pressure is far from sufficient to justify revising the belief. Put differently, that fact provides an epistemic doxastic reason that insufficiently supports Jane giving up her belief. Yet this insufficient reason is sufficient to induce a certain degree of epistemic worry. To mark this type of epistemic worry, we can call it a 'minor epistemic worry.' That is a sort of epistemic worry for a belief which is caused by an insufficient epistemic doxastic reason against that belief, and therefore, should not be eased by revising the relevant belief. Hence, in coming to know the news, Jane can rightly sense a minor epistemic worry with regard to her belief.

Let's consider the second question: Can Jane justify calling the library by pointing to this piece of news? I aim to argue that she can, and that the reason that supports her action is an epistemic practical reason. To this end, I need to show that what Jane comes to know about the incident has an epistemic normative force that supports a course of action.

A common strategy to argue that a consideration has normative epistemic force is to point out its connection with epistemic blame.¹³ One can fittingly epistemically blame only if the subject of the blame is an instance of a normative failing (Boult 2021c, 518). Any normative failing is a matter of a poor response to a reason, and therefore, showing that a particular belief or action deserves epistemic blame is to indicate that there is a relevant epistemic reason which counts against it. This strategy cannot help us, however, since we seek to prove that a commendatory reason supports calling the library and a commendatory reason can never ground for any kind of blame.

Instead, we can rely on the notion of epistemic praise. A consideration that can provide the ground for legitimate epistemic praise without ever licensing any sort of epistemic blame deserves to be called an epistemic reason. Our practice of blaming is not the only tool we have to detect reasons. Another tool of the same importance is our practice of praising people's responses to reasons. Therefore, I suggest the following criterion for epistemic reasons:

The hallmark of an epistemic reason: *r* is an epistemic reason if and only if one's response to it can provide the ground for assigning or blocking epistemic blame or praise.

To show that Jane's reason to call the library is an epistemic reason, I need to show that calling the library is epistemically praise- or blameworthy. Given the alleged commendatory nature of Jane's reason, showing its praiseworthiness is the only option.

Though one can find an extensive discussion of epistemic blame in the literature, the concept of epistemic praise is largely underexplored. However, the available accounts of epistemic blame can act as a source of inspiration for our discussion. Building upon Scanlon's view of moral blame, Cameron Boult has recently developed an account of epistemic blame in terms of relationship modification. The basic idea is that there is a distinctively epistemic relationship between agents: they are "a paradigm source of testimonial knowledge" for each other. The bedrock of such a relationship is epistemic trust—i.e., to take another agent as a source of testimonial knowledge involves having a certain degree of trust in them. Thus, trust is a constitutive element of an intersubjective epistemic relationship. Epistemically blaming an agent involves modifying this specific epistemic relationship with her. Based on Boult's account, we can maintain that a response to an epistemic reason is epistemically blameworthy if and only if it makes the agent less trustworthy as a believer.

¹³Sebastian Schmidt (2021) uses this strategy to argue that evidence has a normative force. Carolina Flores and Elise Woodard (2023) use the same strategy to argue that we have an epistemic duty to gather evidence in certain circumstances. They all adopt the account of blame developed by Cameron Boult (2021a, 2021c). For an overview of the literature on epistemic blame see Boult (2021b).

But using this strategy is not limited to the discussions in epistemology. For instance, Robbie Kubala uses a similar strategy to argue for aesthetic obligations (2020).

This can give us a clue about how to understand the notion of epistemic praise.¹⁴ Epistemically praising another agent also involves modifying our epistemic relationship with her. That is, her credibility as a source of testimonial knowledge increases for us. Therefore, a response to an epistemic reason is epistemically praiseworthy if and only if it makes the agent more trustworthy as a believer.

As mentioned above, the news about the forgery induces a minor epistemic worry in Jane about the veracity of her belief. Calling the library is an epistemic practice whose aim is to ease this minor epistemic worry. The question is the following: Does such a response (i.e., calling the library) make Jane more trustworthy as a believer? The answer seems to be positive. By showing sensitivity to even minor epistemic worries, the agent makes herself a better source of testimonial knowledge and, ipso facto, more trustworthy. But for Jane's action to be epistemically praiseworthy, calling the library must have an epistemically normative significance. It must be in response to an epistemic reason. This shows that Jane must have an epistemic reason to call the library. Moreover, such a reason favors engaging in a course of action and therefore it must be an epistemic practical reason.

Still, even in the absence of countervailing reasons, Jane is not required to call the library. In response to a purported criticism such as "Why didn't you call the library then?" she can simply say "I lost my interest in the topic," or "I don't care about the issue anymore." This response does not make Jane less trustworthy or blameworthy as a believer. It is not the case that if calling the library makes Jane more trustworthy, not calling will make her less trustworthy. As far as evidence is concerned, her belief accords with the available evidence, and this is what we can epistemically expect from a subject. Any higher expectation needs further justification based on a nonepistemic consideration—e.g., the institutional role of Jane as a researcher. If we cannot legitimately expect Jane to call the library, we cannot reduce or suspend our trust when she decides not to call. Therefore, Jane's reason to call the library is a commendatory reason.

Note that my claim is not that the reason that puts pressure on Jane's belief is identical to the reason that supports her calling the library. Rather, I claim that a single consideration provides Jane with two reasons: an epistemic doxastic reason and an epistemic practical reason. But these two reasons are not independent. The piece of news about the library cannot give Jane an epistemic practical reason without providing an epistemic doxastic reason that causes some minor epistemic worry.¹⁵

At this stage, there is a crucial question that we need to address. Do all facts that provide an epistemic doxastic reason against our beliefs also provide us with an epistemic practical reason? After all, such considerations often (or perhaps always) indicate a probability of error. In the above scenario, after calling the library, Jane realizes that more than ninety percent of their historical documents in the library have been replaced by forged documents. Does this second piece of information give her an epistemic practical reason to engage in another epistemic practice? To answer this question, I need to introduce the concept of an error probabilifier.

An 'error probabilifier' is an event that under a certain description indicates the probability of error for a given doxastic state. As an event, any error probabilifier is amenable to a multitude of

¹⁴To the best of my knowledge, Boulton does not discuss the notion of epistemic praise at all. In a passing comment, Flores and Woodard signal that they also take epistemic praiseworthiness to be related to trust. But they do not elaborate on the nature of this relationship or how epistemic praise is connected to epistemic reasons. Instead, they mainly focus on the notion of epistemic blame and its connection to epistemic duties (Flores and Woodard 2023).

¹⁵Pamela Hieronymi famously argued that evidence cannot count in favor of an action, and that a practical reason cannot count in favor of a belief (2005). They simply concern different questions. My position in this paper is compatible with this thesis. In my terminology, her position is that an epistemic doxastic reason can never be a practical reason. Some might believe that her thesis implies that there cannot be epistemic reasons for actions. But this follows only if we identify epistemic reasons with evidence. My argument shows that some practical reasons are associated with epistemic praise. Therefore, they deserve to be categorized as epistemic reasons, not evidence. Thus, her idea that evidence cannot count in favor of actions is compatible with my view. I should add that Hieronymi accepts that a single consideration can provide us with both evidence and practical reasons (439).

correct descriptions.¹⁶ It can indicate the probability of error under certain descriptions, but not others. However, we take an event to be an error probabilifier for a doxastic state as long as there is at least a single description under which the event indicates the probability of error. We can call such a description an ‘error-indicating description.’

Each error-indicating description provides an epistemic doxastic reason that count against the relevant doxastic state. However, not all error-indicating descriptions of the same error probabilifier are on par. Sometimes, the probability of error indicated by them is not high enough to support giving up the relevant belief. Still, it is enough to prompt a minor epistemic worry. In such cases, the subject finds an epistemic practical reason since the engendered minor epistemic worries should not be eased by revising the relevant belief.

In **Library**, Jane encounters two error-indicating descriptions of the same event—i.e., the replacement of the historical documents with forged and misleading ones. The first description (i.e., a forgery has happened in one of the eight hundred libraries in the city) does not give Jane sufficient reason to give up her beliefs. Based on this error-indicating description, Jane cannot know about the scope and the time of the replacement. Perhaps the forgery was limited to a few documents, or it happened very recently, after she finished her research, or in other libraries of the city. But the description alludes to an error probabilifier whose presence prompts a minor epistemic worry.

After Jane calls the library, she finds a more detailed error-indicating description of the event (i.e., more than ninety percent of the historical documents in the library have been replaced by forged documents). This time the probability of error is so high that she needs to give up her belief. Here, what she comes to know does not give her any epistemic practical reason since it does not prompt a minor epistemic worry. Given that she needs to give up her belief, there is nothing left to be epistemically worried about. The subject may have an epistemic practical reason as long as she can rationally retain the belief that is the subject of epistemic worry.

Thus, what distinguishes the first error-indicating description from the second one is that the first one prompts a minor epistemic worry. *It causes epistemic worry since it alludes to an error probabilifier that might under a more precise description give the subject a reason to give up her belief.* This is why the first error-indicating description provides Jane with an epistemic practical reason.

Before closing this section, I would like to briefly discuss the kinds of epistemic activity that one can justify by citing an error-indicating description. In **Library**, Jane’s epistemic practical reason intuitively supports calling the library. The rationale behind this action is straightforward: she needs to acquire a more precise description of the error probabilifier. Her attempt to gain a more precise description of the error probabilifier is itself an epistemic activity, or more precisely, an inquiry. The aim of such an inquiry is to find out whether the impact of the error probabilifier was sufficient to provide her with a defeater for her belief. But this is not the only action that Jane can justify by referring to the first error-indicating description.

In many cases, we might not have the means to find a more precise description of the error probabilifier. We sense the probability of an error, and we find it epistemically worrisome without having a way to find out whether the impact of the error probabilifier is significant enough to undercut our belief. Still, in being epistemically worried about our doxastic state, we do have a practical reason to minimize or exclude the probability of error—for instance, by gathering more evidence supporting our belief. To see this, assume that in **Library**, the receptionist refuses to give any information about the event to Jane. In such a case, Jane finds a reason to gather more first-order evidence supporting her belief regarding Catherine the Great—e.g., by looking into the historical documents in other libraries. After all, if she can back up her belief by finding similar

¹⁶The idea that an event can be described using different descriptions is a famous Anscombe-Davidsonian insight (see Anscombe 1979; Davidson 1969).

documents in other libraries, it ceases to matter whether the documents in the first library were forged or not. As a result, the minor epistemic worry will be eased.

3. Addressing the demarcation worry

At this stage, we are equipped with the required conceptual framework to address the challenge of irrelevant influences. If successful, my suggested analysis of irrelevant influences answers two important questions. First, it helps us answer how to demarcate three types of scenarios—i.e., the doomed cases, the gray cases, and the innocuous cases. Second, and more importantly, it helps us answer what the proper response in facing an irrelevant influence is. In this section, I will focus on the first question. In the next section, I will present my answer to the second one.

It is best to start our discussion by considering the connection between the probability of error and the challenge of irrelevant influences. This will help us relate the challenge at issue with the notion of error probabilifiers and the framework developed in the previous section. Some epistemologists have already pointed out that there is such a connection. For instance, both Katia Vavova and George Sher argue that this connection is what differentiates the challenge of irrelevant influences from skeptical challenges.¹⁷ In a typical skeptical challenge, the skeptic provides us with a scenario that merely indicates the possibility of error, not its probability. Then, he shows that we cannot exclude this possibility, and therefore, we cannot know that we are not, say, a brain in a vat. By contrast, finding an irrelevant factor indicates a real probability of error, not just its possibility. This claim is backed up by some empirical findings that demonstrate the biasing influence of irrelevant factors, such as desire or genetics.¹⁸ But the significance of the correlation between the probability of error and the challenge of irrelevant influences goes beyond showing its difference from skeptical challenges. It can also explain some of the more puzzling features of irrelevant influences.

To see this, we need to understand the challenge of irrelevant influences and its related concepts, such as an IF consideration, from within the framework developed in the previous section. My suggestion is that we take the presence of an irrelevant factor to be an error probabilifier—an event that, under a description, indicates the probability of error for a given doxastic state. An epistemically significant IF consideration is an error-indicating description of such an event that creates some pressure to revise the impacted belief. As explained in the previous section, an error-indicating description can provide us with different types of epistemic reasons. Ipso facto, an epistemically significant IF consideration does not necessarily provide us with just an epistemic doxastic reason. In some cases, it can also give us an epistemic practical reason to engage in epistemic activities.

The above move helps us correct a common misconception in the literature. There is a consensus in the literature on what I called the doxastic interpretation of IF considerations according to which an epistemically significant IF consideration is supposed to provide a defeating epistemic doxastic reason, more precisely a piece of defeating higher-order evidence. Here, we can see that this might not always be the case. Sometimes, an epistemically significant IF consideration can cause epistemic worry without providing us with a sufficient reason to revise the impacted belief. But we should not hasten to conclude that the IF consideration in such a case does not provide us with any other reason. Based on the framework developed in the previous section, it can provide us with an epistemic practical reason that renders engaging in a course of action epistemically choiceworthy. Therefore, *the epistemic significance of an IF consideration is not just doxastic, it can also provide us with an epistemic practical reason.* This is why we should abandon the doxastic interpretation in

¹⁷See Vavova (2018) and Sher (2001, 66–67); Some have appealed to the same distinction to explain the difference between skeptical challenges and cases of peer disagreement (Rattan 2014, 2018).

¹⁸See Avnur and Scott-Kakures (2015) and Munro (2020).

favor of what I called the hybrid interpretation of IF considerations according to which an epistemically significant IF consideration either provides us with a defeating epistemic doxastic reason or an epistemic practical reason.

We are now in a position to address the demarcation worry introduced in the first section. The question was how we can differentiate three types of scenarios—i.e., the doomed cases, the gray cases, and the innocuous cases. I believe that what determines which of the three types a scenario belongs to is not the irrelevant factor itself, but the description of it available to the subject. Thus, what the subject's proper response should be can vary depending on what description of the error probabilifier is available to him. Therefore, it is possible to construct, for instance, a doomed case and an innocuous case through two different descriptions of the same irrelevant factor.

This is an observation that often gets lost in the literature on irrelevant influences. The problem, as it is often framed, concerns whether learning about the irrelevant factor can defeat the subject's justification. Yet an oft-overlooked point is that we can learn about the same irrelevant factor under different descriptions. The proper response, then, can vary depending on which error-indicating description we come to know.

We can start with doomed cases in which becoming aware of an irrelevant factor provides us with a defeater. Consider a scenario in which I realize that I was under the influence of a very competent hypnotist when I formed the judgment that Ayer's argument in a particular passage is plausible. The relevant IF consideration, in this case, is the following: my belief in the plausibility of Ayer's argument is formed under the influence of a hypnotist. This consideration provides me with a defeating epistemic doxastic reason. The reason is that the error-indicating description shows that the probability of error is high enough to require me to revise the impacted belief. This explains why in some cases of irrelevant influences, we believe that the subject needs to give up his belief.

By contrast, in a gray case, the probability of error is not high enough to provide the subject with a defeater, though it prompts minor epistemic worry. In the example of *The Medical Historian*, Sarah points out to Silvia that she finds her findings about the Soviets conclusive because of her upbringing in a communist household. The relevant IF consideration, in this case, points to an error probabilifier—i.e., Silvia's upbringing in a communist household. Though such a consideration puts some limited pressure on Silvia's belief and prompts minor epistemic worry, it is not strong enough to rationally compel her to give up her belief. Still, it provides a practical epistemic reason for engaging in the relevant epistemic activities to assuage the minor epistemic worry. If that's true, *the gray cases vindicate the claim that the epistemic significance of IF considerations goes beyond their status as higher-order defeaters*. No view can fully capture in which sense IF considerations are or are not epistemically significant without taking this point into consideration. But, to the best of my knowledge, the available views in the literature have overlooked this point.

Here, one might rightly ask which sort of epistemic activity can fulfill this task. This is a topic that we will discuss in detail in the next section. For this section, it suffices to see that the relevant IF consideration provides the subject with an epistemic practical reason. Given our discussion in the previous section, such a reason does not require but only commends engaging in a course of action.¹⁹

¹⁹In this paper, I have exclusively focused on qualitative doxastic attitudes—viz., belief, disbelief, and suspension. However, one might wonder whether the IF consideration in a gray case gives the subject any reason to revise her level of confidence. A satisfactory response to this question would require a comprehensive discussion of the relationship between belief and credence that goes beyond the scope of this paper. However, I can explain what the response might be. I am more inclined toward the so-called belief-credence dualism, according to which both belief and credence are fundamental. If this view is correct, an IF consideration can provide the subject with three distinct reasons: a sufficient reason to reduce her level of confidence, an insufficient reason to revise her belief (which causes a minor epistemic worry), and an epistemic practical reason. For a defense of dualism, see Jackson (2018, 2019) and Friedman (2019). However, this is not necessarily true if you have a different understanding of the relationship between belief and credence. For instance, if you believe that belief is credence 1 (namely,

In an innocuous case, we usually have a factor that under another description can indicate the probability of error for a given doxastic state. However, the description available to the subject indicates negligible or zero probability of error. Consider the innocuous case discussed in the first session, in which my child's belief about the litmus test has been formed partly due to his desire for a reward. In this case, his desire has played a role in the formation of his belief. But our description of the case shows that there is no room for any kind of epistemic worry. As a result, the available description of the factor provides neither an epistemic doxastic reason nor an epistemic practical reason.

To wrap up, the normative significance of the available description of the irrelevant factor in each type of scenario is different. In a doomed case, it provides the subject with a defeating epistemic doxastic reason. In a gray case, it gives the subject a commendatory epistemic practical reason to engage in a course of action. Finally, in an innocuous case, it does not provide the subject with any relevant doxastic or epistemic practical reason.

So far, I have explained how to demarcate the three types of scenarios that involve the presence of an irrelevant factor. We still need to clarify how a subject can respond upon identifying an irrelevant factor in the pedigree of her beliefs. Answering this question will be the subject of the next section.

4. The challenge of irrelevant influences

If my arguments up to this point have been convincing, we must admit that there is no unique answer to the challenge of irrelevant influences. In my adopted formulation, the challenge of irrelevant influences is to answer the following question: How can we respond if we identify an irrelevant factor in the pedigree of our beliefs? We always identify irrelevant factors under the guise of a description. As we saw in the previous section, the normative forces of these descriptions vary in different types of cases. As a result, our answer to the challenge of irrelevant influences cannot be the same for all three types of cases.

The proper answers to the challenge of irrelevant influences in an innocuous case and a doomed case are more or less obvious. The description of the irrelevant factor in an innocuous case does not provide the subject with any epistemic practical or epistemic doxastic reason. Thus, it does not prompt any epistemic worry. The subject can simply dismiss the presence of the irrelevant factor, though she might realize that under a different error-indicating description the same factor can instigate epistemic worry for her belief. By contrast, in a doomed case, the IF consideration provides the subject with a defeating epistemic doxastic reason. As a result, the subject is rationally compelled to give up the impacted belief, and the proper response is to do so. As argued in the second section, in such cases, the subject does not find a further epistemic practical reason to engage in any epistemic activity.

However, it takes more effort to determine what the proper response is in a gray case. I claimed that in such a case, the IF consideration provides epistemic practical reasons to engage in certain epistemic activities. Given the commendatory nature of such reasons, the subject has both the option to dismiss the IF consideration and the option to engage in the commended epistemic activity. But this only partially answers how the agent can respond in a gray case. To fully address the question, we need to determine which epistemic activities are commended by an IF consideration in a gray case.

To answer this question, I will focus on the example of **The Medical Historian** depicted in the first section. In this example, Silvia is a medical historian who believes her findings point to a great success achieved by the Soviets in controlling a pandemic. However, Sarah puts forward an IF consideration describing an irrelevant factor that is present in the formation of Silvia's belief

maximal credence), then an IF consideration in a gray case does not provide the subject with sufficient reason to reduce her level of confidence. For a defense of such a view on credence, see Wedgwood (2012) and Dodd (2016).

—i.e., her upbringing in a communist household which might have instilled in her an unconscious affinity for communists. As explained in the first section, this is a gray case, since becoming aware of the irrelevant factor under the aforementioned description does not provide Silvia with a defeater. Yet it prompts a minor epistemic worry, and therefore provides her with epistemic practical reasons.

In the second section, I claimed that an error-indicating description can provide the subject with a reason to engage in an inquiry to find a better description of the error probabilifier (e.g., calling the library in **Library**). However, in cases where the subject does not have the necessary means to find a more precise description, she will find a reason to engage in other epistemic activities aiming to assuage the minor epistemic worry (e.g., in **Library**, if the receptionist refuses to answer Jane's inquiry, she finds a reason to look into the historical documents in other libraries). Such activities assuage the minor epistemic worry by excluding or minimizing the probability of error caused by the error-indicating description.

This framework gives us a clue as to what the relevant epistemic activities are in a gray case like **The Medical Historian**. First, Silvia finds an epistemic practical reason to seek a more precise description of the error probabilifier. If she fails to find a more precise description, she has a reason to engage in other epistemic activities aiming to assuage the minor epistemic worry by minimizing or excluding the probability of error. I will examine each of these possibilities in turn.

4.a One reason to find a more precise description of the irrelevant factor

In a gray case, one finds a reason to gather more information about the impact of the irrelevant factor on one's belief. But someone can get more information about the irrelevant factor without being able to assuage her epistemic worry. Thus, it is better to maintain that the IF consideration in a gray case provides the subject with a reason to find a *sufficiently precise description* of the error probabilifier (i.e., the presence of the irrelevant factor). A description is sufficiently precise if and only if it either shows that the presence of the irrelevant factor poses no epistemic pressure for the revision of the impacted belief, or it gives the subject a sufficient reason to give up her belief. In other words, an IF consideration commends finding a more precise description that turns the case into a doomed or an innocuous case.

This is also true of Silvia in **The Medical Historian**: she finds a reason to seek a sufficiently precise description of the irrelevant factor. For instance, she could take psychological tests that illuminate the impact of her upbringing on her judgments about issues related to communists. Alternatively, she could seek help from a psychoanalyst to investigate his father's influence on her biases. Such practices might help her to find a more precise description of the impact of her upbringing on her evidence processing.

There is, however, a genuine obstacle in Silvia's way. Though there are cases in which the agent in question can easily find an answer to such an inquiry, in many cases, including hers, there is no clear-cut path to acquire a sufficiently precise description of the error probabilifier. Clearly, her upbringing could have biased her in favor of communists, and therefore it can be an error probabilifier. But determining the impact of the upbringing in each individual case is not an easy task. Thus, it is difficult for Silvia to determine whether the impact was or was not significant enough to distort her evidence processing. As a result, she might not be able to find a sufficiently precise description of the error probabilifier.

What is the path forward? In many cases, a single consideration can provide us with a reason to engage in a course of action, and when we fail to achieve the desired goal, it provides us with a reason to engage in a different course of action. The fact that my friend's chest is bleeding primarily gives me a reason to call 911 and wait until a professional health practitioner helps him. But if I realize that I cannot contact 911, the same fact (*that his chest is bleeding*) gives me a reason to try to stop the bleeding myself. In **The Medical Historian**, Sylvia's primary reason is to seek a sufficiently precise description of the error probabilifier. But if she does not have the necessary means to do so, she will

find a reason to engage in other epistemic practices to minimize the probability of error. To the discussion of these epistemic practices, we now return.

4.b Reasons to minimize the probability of error

In this subsection, I will consider two other types of epistemic activities that can help Silvia to assuage her minor epistemic worry. First, she could gather more first-order evidence supporting her impacted belief. Second, she could gather second-order evidence indicating that first-order evidence she already possessed supports the impacted belief. I will show how each of these activities can help Silvia to exclude or minimize the probability of error, and therefore, assuage her minor epistemic worry.

4.b.1 One reason to gather more first-order evidence

So far, I have mostly written as if we calculate the probability of error based merely on the error-indicating description. But surely this is not a fully accurate picture. Facing the same error-indicating description, two subjects can rightly sense different probabilities of error due to the difference in the strength of their first-order evidence. The stronger the evidence, the weaker the probability of error indicated by an error-indicating description. To see this, consider the following scenario:

Conservative Family: Dan is brought up in an ultraconservative family. Due to his father's abusive behavior, Dan develops a mild bias against his father's system of beliefs. In particular, he becomes biased against his antiscientific beliefs, in particular his affinity for the geocentric model of the universe. Though this bias is not strong enough to completely blind him to countervailing evidence against heliocentrism, it has impacted his ability to judge. After finishing high school, Dan goes to college and studies science. Finally, he pursues studies in astronomy and completes his PhD, becoming an expert in the solar system.

In the above scenario, Dan's bias against his father's system of belief plays a role in the formation of his belief that the sun does not orbit the earth. Yet it is not the case that he has become completely blind to countervailing evidence. Is it plausible to criticize Dan by pointing out that he believes in heliocentrism due to his hostility to his father's system of belief? Given all he knows regarding the solar system, it hardly makes sense for Dan to revise or even feel any epistemic worry in response to such criticism. The criticism is misplaced and, therefore, despite the biasing effect of Dan's feelings, **Conservative Family** is an innocuous case. Why is that?

The reason is that the impact of irrelevant influences, like any other error probabilifier, has certain limits. It is not that the presence of a biasing effect can always completely impair our ability to reliably form beliefs. Given all evidence available to Dan, the probability of error is indeed slim, and this is why **Conservative Family** is an innocuous case.²⁰

The above example teaches us an important lesson regarding irrelevant factors. By gathering better and more reliable evidence, we can be in a better epistemic position with regard to an irrelevant factor. If the strength of evidence possessed by the subject is such that the probability of error becomes negligible, the IF consideration does not prompt any epistemic worry.

Here, we can see that finding out about the presence of an irrelevant factor gives the subject a reason to seek more first-order evidence. This is an epistemic activity, more precisely an inquiry, that can ease her epistemic worry. For instance, in **The Medical Historian**, facing Sarah's IF criticism gives Silvia a practical reason to engage in an inquiry to find more evidence supporting her

²⁰Note that if Dan's feelings were so strong that he had become completely blind to countervailing evidence, we could say that even after getting his PhD, he lacks doxastic justification. The reason is that in such a case, his belief could not be based on the evidence he possesses.

belief. By finding better first-order evidence, similar to Dan in **Conservative Family**, the IF consideration stops prompting any epistemic worry in her.

4.b.2 *One reason to gather more second-order evidence*

Finally, there are cases in which we cannot pursue either of the above epistemic activities. Imagine that Silvia cannot find more documents or historical witnesses supporting her claim. Further, assume that she has no means to find a sufficiently precise description of the error probabilifier. In such a case, I believe she still has a reason to engage in a different epistemic activity. She can try to gather second-order evidence showing that the first-order evidence she already possessed supports her belief.

There is a limit to what such second-order evidence can show. It does not necessarily show that Silvia's upbringing has not negatively influenced her evidence-processing ability. Yet it can show that she already had the relevant propositional justification for her belief. Arguably, this leaves room for her not having doxastic justification despite having propositional justification. To fully vindicate her evidence-processing ability, she needs to engage in the first kind of inquiry described above. Still, finding second-order evidence showing that her already-held first-order evidence supports her belief can minimize the possibility of her belief being erroneous.

Similar to the two previous epistemic activities, this activity is also a sort of inquiry. The best example of such inquiry can be the act of convincing a neutral person—i.e., a person who has not been influenced by the same irrelevant factor—by citing the reasons for our belief. By presenting our evidence to a neutral subject, and trying to convince him, we engage in an inquiry to see whether our body of evidence is sufficient to support the same belief without the interference of the irrelevant factor. For instance, Silvia can see whether subjects who do not share her upbringing also find her conclusion plausible after studying her first-order evidence. This is why we find our success in convincing others to be epistemically assuring. It provides us with second-order evidence showing that we already held propositional justification for our belief.

Now we have the full picture of my suggested answer to the challenge of irrelevant influences. My aim was to show that we fail to fully appreciate the epistemic significance of IF considerations if we focus merely on epistemic doxastic reasons. In some cases, such considerations commend our engaging in a course of action. Consequently, no complete answer to the challenge of irrelevant influences can dismiss epistemic practical reasons.

5. Conclusion

In this paper, I suggested a way to address the challenge of irrelevant influences—i.e., to answer the question: How can we respond if we identify an epistemically worrying irrelevant factor in the pedigree of our beliefs? In a doomed case, the agent needs to revise his belief. In a gray case, the subject finds three commendatory reasons to engage in three epistemic activities. In an innocuous case, he can simply dismiss the presence of the irrelevant factor. If successful, my arguments in this paper show that a proper response to the challenge of irrelevant influences cannot focus merely on doxastic states and epistemic doxastic reasons. It also needs to include what I called epistemic practical reasons and epistemic activities.

Acknowledgments. I am grateful to Gurpreet Rattan for his insightful feedback, which helped me develop my ideas and improve the quality of my writing. I would also like to thank Hadi Safaei, Daniel Z. Korman, Eliran Haziza, Nate Charlow, Nathan Howard, Sergio Tenenbaum, and Zain Raza for their helpful comments on earlier drafts of the paper. Their feedback helped me to clarify my thoughts and improve my argument.

Seyed Mohammad Yarandi is an assistant professor, Teaching Stream in the Department of Philosophy at the University of Toronto Scarborough. He works primarily on ethics, epistemology, and Islamic philosophy.

References

- Anscombe, G. E. M. 1979. "Under a Description." *Noûs* 13 (2): 219–33.
- Avnur, Yuval, and Dion Scott-Kakures. 2015. "How Irrelevant Influences Bias Belief." *Philosophical Perspectives* 29: 7–39.
- Booth, Anthony Robert. 2006. "Can There Be Epistemic Reasons for Action?" *Grazer Philosophische Studien* 731 (1): 133–44.
- Boult, Cameron. 2021a. "The Significance of Epistemic Blame." *Erkenntnis* 88: 1–22.
- Boult, Cameron. 2021b. "Epistemic Blame." *Philosophy Compass* 16 (8): e12762.
- Boult, Cameron. 2021c. "There Is a Distinctively Epistemic Kind of Blame." *Philosophy and Phenomenological Research* 103 (3): 518–34.
- Dancy, Jonathan. 2004. "Enticing Reasons." In *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, edited by R. Jay Wallace, Phillip Pettit, Samuel Scheffler, and Michael Smith, 91–118. Oxford: Clarendon Press.
- Davidson, Donald. 1969. "The Individuation of Events." In *Essays in Honor of Carl G. Hempel*, edited by Nicholas Rescher on behalf of the editorial committee, 216–34. Dordrecht, Nether.: Springer.
- Dodd, Dylan. 2016. "Belief and Certainty." *Synthese* 194 (11): 4597–621.
- Dworkin, Ronald. 1996. "Objectivity and Truth: You'd Better Believe It." *Philosophy & Public Affairs* 25 (2): 87–139.
- Dworkin, Ronald. 2011. *Justice for Hedgehogs*. Cambridge, MA: Harvard University Press.
- Elga, Adam. 2008. "Lucky to Be Rational." Paper presented at the Bellingham Summer Philosophy Conference, June.
- Feldman, Richard. 2000. "The Ethics of Belief." *Philosophy and Phenomenological Research* 60 (3): 667–95.
- Feldman, Richard. 2002. "Epistemological Duties." In *The Oxford Handbook of Epistemology*, edited by Paul K. Moser. New York: Oxford University Press.
- Flores, Carolina, and Elise Woodard. 2023. "Epistemic Norms on Evidence-Gathering." *Philosophical Studies* 180: 2547–71.
- Friedman, Jane. 2019. "Inquiry and Belief." *Noûs* 53 (2): 296–315.
- Gert, Joshua. 2004. *Brute Rationality*. Cambridge: Cambridge University Press.
- Hedden, Brian. 2015. *Reasons without persons: Rationality, identity, and time*. Oxford University Press.
- Hieronymi, Pamela. 2005. "The Wrong Kind of Reason." *Journal of Philosophy* 102 (9): 437–57.
- Horgan, Terry, and Mark Timmons. 2010. "Untying a Knot from the Inside Out: Reflections on the "Paradox" of Supererogation." *Social Philosophy and Policy* 27 (2): 29–63.
- Jackson, Elizabeth. 2018. "A Defense of Belief-Credence Dualism." In *The Proceedings of the of the Brazilian Society of Analytic Philosophy*, edited by João Luis Pereira Ourique, 77–78. Pelotas, Brazil.
- Jackson, Elizabeth. 2019. "How Belief-Credence Dualism Explains Away Pragmatic Encroachment." *The Philosophical Quarterly* 69 (276): 511–33.
- Kelly, Thomas. 2003. "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 66 (3): 612–40.
- Kubala, Robbie. 2020. "Aesthetic Obligations." *Philosophy Compass* 15 (12): e12712.
- Lord, Errol. 2018. "Epistemic Reasons, Evidence, and Defeaters." In *The Oxford Handbook of Reasons and Normativity*, edited by Daniel Star, 601–28. Oxford: Oxford University Press.
- Little, Margaret Olivia, and Coleen Macnamara. 2020. "Non-Requiring Reasons." In *The Routledge Handbook of Practical Reason*, edited by Ruth Chang and Kurt Sylvan, 393–404. London: Routledge.
- Little, Margaret Olivia, and Coleen Macnamara. 2022. "The Latitude Preserving Nature of Commendatory Reasons." *Oxford Studies in Metaethics* 17: 54–76.
- Mogensen, Andreas L. 2015. "Contingency Anxiety and the Epistemology of Disagreement." *Pacific Philosophical Quarterly* 97 (4): 590–611.
- Munroe, Wade. 2020. "The Challenge of Heritability: Genetic Determinants of Beliefs and Their Implications." *Inquiry* 63 (8): 831–74.
- Nelson, Mark T. 2010. "We Have No Positive Epistemic Duties." *Mind* 119 (473): 83–102.
- Peels, Rik. 2016. *Responsible Belief: A Theory in Ethics and Epistemology*. New York: Oxford University Press.
- Rattan, Gurpreet. 2014. "Disagreement and the First-Person Perspective." *Analytic Philosophy* 55 (1): 31–53.
- Rattan, Gurpreet. 2018. "Disagreement and Conceptual Understanding." *Theoria* 84 (2): 179–210.
- Scanlon, Thomas M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Scanlon, Thomas M. 2014. *Being Realistic about Reasons*. Oxford: Oxford University Press.
- Schmidt, Sebastian. 2021. "Epistemic Blame and the Normativity of Evidence." *Erkenntnis*: 1–24.
- Schoenfield, Mark. 2014. "Permission to Believe: Why Permissivism Is True and What It Tells Us about Irrelevant Influences on Belief." *Noûs* 48 (2): 193–218.
- Schroeder, Mark. 2021. *Reasons First*. Oxford: Oxford University Press.
- Sher, George. 2001. "But I Could Be Wrong." *Social Philosophy and Policy* 18 (2), 64–78.
- Vavova, Katia. 2018. "Irrelevant Influences." *Philosophy and Phenomenological Research* 96 (1): 134–52.
- Wedgwood, Ralph. 2012. "Outright Belief." *Dialectica* 66: 309–29.
- White, Roger. 2010. "You Just Believe That Because ..." *Philosophical Perspectives* 24 (1): 573–615.