

## *How Supererogation Can Save Intrapersonal Permissivism*

*Han Li*

### **Abstract**

Rationality is intrapersonally permissive just in case there are multiple doxastic states that one agent may be rational in holding at a given time, given some body of evidence. One way for intrapersonal permissivism to be true is if there are epistemic supererogatory beliefs – beliefs that go beyond the call of epistemic duty. Despite this, there has been almost no discussion of epistemic supererogation in the permissivism literature. This paper shows that this is a mistake. It does this by arguing that the most popular ways of responding to one of the major obstacles to any intrapersonally permissive all fall prey to the same problem. This problem is most naturally solved by positing a category of epistemically supererogatory belief. So intrapersonal epistemic permissivists should embrace epistemic supererogation.

### *1. Introduction*

Rationality constrains the beliefs we are allowed to have – this much is obvious. In any situation, there are certain beliefs that are simply irrational to hold. What is less obvious is *how much* rationality constrains our beliefs. How strict is rationality? Is the rational agent chained by the dictates of reason, such that her opinions are shaped only by the contingencies of her situation? Many writers intuitively think (and perhaps many more hope), that this is not true. That is, maybe even the most zealous follower of reason is allowed some latitude in her beliefs. Maybe rationality allows us some leeway in what we ultimately believe. Perhaps rationality is “permissive.”

Here is one way that rationality might be permissive. Even if there is always one unique belief that is maximally rational to hold in any given situation, perhaps agents are not irrational if they fail to hold this belief. Sometimes, there are beliefs that are rationally permissible to hold even if they are not maximally rational. Maybe rationality sometimes cuts us slack, letting it go even if we don't always do the optimal thing. For rationality, sometimes good enough is good

enough. What I am suggesting is that there may be epistemically supererogatory beliefs – beliefs that go beyond the call of epistemic duty. This opens up the conceptual space for beliefs that might meet epistemic duty – are rationally permissible to form – even though there are better beliefs available. There is, of course, a large and varied literature on moral supererogation. But there has been very little written on the epistemic counterpart.<sup>1</sup>

We should find this surprising. At the very least, there are many seeming parallels between ethics and epistemology, the two major normative philosophical disciplines. Investigating these parallels is a good strategy to make progress in both fields, even if, on the final analysis, the seeming parallels turn out to be misleading.<sup>2</sup> This in itself makes it surprising that epistemic supererogation is rarely discussed. It is even more surprising that discussions of epistemic permissivism rarely talk about supererogation, even though epistemic supererogation is clearly one way for rationality to be permissive.<sup>3</sup> In this paper, I will argue that this is a mistake. Indeed, defenders of a certain kind of permissivism should embrace epistemic supererogation.

Before we begin, some clarifying remarks are in order. Permissivism is traditionally conceived of as the denial of a thesis known as “uniqueness.” In this paper, we will follow this tradition. Here is a fairly orthodox formulation of the thesis, which is more than adequate for our purposes:

---

<sup>1</sup> One exception is Hedberg (2014), although that paper is about epistemically supererogatory actions (such as gathering additional evidence or double checking past evidence), whereas I am interested in epistemically supererogatory doxastic states. As we will see, the theory of epistemic supererogation I favor has much in common with Hedberg’s general approach.

<sup>2</sup> This point is emphasized by Berker (2013) and Hedberg (2014).

<sup>3</sup> Though, as we will see, some proposals can be interpreted as forms of epistemic supererogation. Podgorski (2016) perhaps comes closest, when it says that you can be “doing better than you need” (§13) when you undergo some non-required epistemic processes.

Uniqueness: Necessarily, for any total body of evidence E, and proposition P, there is at most one doxastic attitude to take toward P that is consistent with being rational and having E.<sup>4</sup>

In addition, most permissivists defend an especially strong version of the view, such that there are situations which are permissive between believing some P and believing  $\neg P$  (as opposed to believing P and suspending judgment on P, for example).<sup>5</sup> In this paper, the term “permissivism” will be referring to the strong version of the view.

Permissivism can be further delineated into two more specific claims, which we will call interpersonal permissivism and intrapersonal permissivism.<sup>6</sup> Roughly, interpersonal permissivism says that two (or more) persons may have different doxastic responses to the same body of evidence and both be rational. Intrapersonal permissivism, on the other hand, says that the *same* person may have one of two (or more) doxastic responses to the same body of evidence and be rational in either response. More specifically, in this paper intrapersonal permissivism will refer to the claim that there can be a point in time where a single agent, with some body of evidence, can form either of two (or more) incompatible doxastic states and be rational no matter which way she goes.<sup>7</sup> In this paper we will be exclusively interested in the prospects for

---

<sup>4</sup> This formulation is slightly modified from Schoenfield (2014, pp. 3). In this paper, I will assume that an agent is rational insofar as she has epistemically justified beliefs. Accordingly, I will use the phrases “rational belief” and “justified belief” interchangeably.

<sup>5</sup> In this paper, I will be talking in terms of all-or-nothing beliefs, although I suspect everything I say can be translated into talk of credences.

<sup>6</sup> This terminology is from Kelly (2013).

<sup>7</sup> The point of using the terminology in this way is to distinguish intrapersonal permissivism from what we might call “possible worlds permissivism,” such that the same person, *in two different possible worlds*, might be rational in having different beliefs in those two worlds even with the same evidence. This view seems weaker than intrapersonal permissivism (in my sense), and is normally held by interpersonal permissivists who do not want to commit themselves to full on intrapersonal permissivism. This is because interpersonal permissivists posit some non-evidential fact about persons that make certain beliefs rational for them. These non-evidential facts can vary across different people, which can also change what is rational for an agent to believe, even with the same body of evidence. But these facts are generally not

intrapersonal permissivism. Thus, the thesis should be understood as limited in this way:  
 intrapersonal permissivists should embrace epistemic supererogation.<sup>8</sup>

It should also be noted that the word “rational” in the definition of uniqueness should not be understood to mean “maximally rational.” This is because we want theories of epistemic supererogation which posit one maximally rational doxastic response to any body of evidence to count as permissive, as long as they also posit some other permitted (though worse) doxastic response. Perhaps because writers typically do not countenance the possibility of epistemic supererogation, the distinction between “maximally rational” and merely “rational” (or “as rational as one is required to be”) is rarely explicitly discussed when attempting to define permissivism.<sup>9</sup> It is possible that after we make the distinction explicit, some defenders of uniqueness will realize that “maximal rationality” is what they were interested in all along. Which means that a theory of epistemic supererogation will be consistent with what these

---

thought of as essential features of any single person – which means that they can also vary across possible worlds. So the same agent can vary with respect to this non-evidential fact across worlds, which means that different doxastic states will be rationalized by the same evidence across worlds. My use of the term is similar to what Podgorski (2016) has called “options-permissivism.”

<sup>8</sup> Though this paper does not argue for the existence of epistemic supererogation for anyone who is not already committed to intrapersonal permissivism, I do think that epistemic supererogation is independently plausible. See Li (2017).

<sup>9</sup> One exception is Christensen (2007) and (2013), who makes his formulation of uniqueness thesis explicitly about maximal rationality. As evidence for the claim that most philosophers are simply not thinking of this distinction, consider the following (incomplete) sampling of how the literature talks about the uniqueness thesis. Some writers such as White (2005), Douven (2009), Cohen (2013), and Schoenfield (2012) simply talk about what is “rational,” while Feldman (2007) defines uniqueness in terms how many propositions the evidence “justifies.” White (2013) and Kelly (2013) define uniqueness in terms of how many doxastic states are “fully rational.” (To my ear, “full rationality” is ambiguous between “maximal rationality” and “as rational as one is required to be.”) Titelbaum and Koepc (ms) present different variations of the uniqueness thesis, with different versions alternately talking about what the evidence “justifies,” “confirms,” or “rationally permits.” Similarly, Ballantyne and Coffman (2011) use both the term “rational” and “justifies” in the same formulation.

philosophers want to call “uniqueness.” Still, enough theorists will agree with our definition of uniqueness for us to reasonably count epistemic supererogation as a form of rational permissivism (as it intuitively is). As we will see, however, there is a sense in which a theory of epistemic supererogation is “closer” to uniqueness than other permissive theories, since it is consistent with a natural line of thought that, *prima facie*, seems to lead to uniqueness. This is another mark in favor of epistemic supererogation.

That being said, not just any theory of epistemic supererogation can serve as a plausible form of permissivism. Consider, for example, a theory which draws the line of rational permissibility by simply encoding everyday judgments of rationality. Clearly, ordinary language allows for deviations from maximum rationality that we nonetheless call “rational” because of the need to make epistemic appraisals of real agents – agents who rarely achieve maximally rational doxastic states. A theory that took these kinds of everyday judgments about rationality as definitive would allow for the existence of epistemic supererogation. Under such a theory, for example, every well-informed mathematician who had a view about some cutting edge research question would probably count as rational, even if they disagreed. Still, the mathematicians who believed the provable true answer would be doing better, and therefore, have supererogatory beliefs.

Though such a theory would allow for epistemic supererogation, it would hardly threaten the truth of uniqueness. Even the most ardent defender of uniqueness will admit that in non-theoretical contexts we talk loosely about rationality. If this is all permissivism comes to, then it would be entirely uncontroversial. Thus, the type of permissivism that such a theory of supererogation represents would hold little theoretical interest. If we are forced to talk about

rationality in the everyday sense, the debate over uniqueness could easily be reinterpreted as a debate about *maximal* rationality.<sup>10</sup>

Fortunately, a theory that relies on everyday ascriptions of rationality and the cognitive capacities of actual agents is not the only option available. To see this, we can examine the analogy with moral supererogation. Many defenders of moral supererogation maintain that there is still a bright line between permissible and impermissible actions. To say that an action is minimally permissible is not simply to say something about its goodness compared to other options – it is not simply saying, for example, that the action is not horribly bad. For even if some permissible actions are not maximally good, performing those actions still discharges one’s moral *duties* – and there is a real moral difference between doing one’s duty and failing to do one’s duty.<sup>11</sup> Thus, we can construct an analogous concept for the epistemic realm, such that rationally permissible doxastic states are epistemically different from the rationally impermissible ones in a theoretically interesting way, even if they are not maximally rational.

In what follows, we will begin with some remarks on the intuitive plausibility of both intrapersonal permissivism and epistemic supererogation. Next, we will consider a general difficulty for intrapersonal permissivism – the “arbitrariness” objection. We will then examine two different types of response to this objection. It will be shown that the two types of response both face the same serious problem as they stand. We will see that positing a category of epistemically supererogatory belief is the natural way to deal with this problem. So it will turn

---

<sup>10</sup> I would like to thank James Fritz for making this point.

<sup>11</sup> Indeed, in his seminal work on moral supererogation, Urmson (1958) explicitly distinguishes between an agent doing her duty (even in contexts where this is extremely difficult, such that ordinary persons would fail to do their duties) and going beyond the call of duty – that is, performing supererogatory actions.

out that different paths to avoid the arbitrariness objection all lead to epistemic supererogation. Finally, we will conclude the discussion by sketching out the skeleton of a positive theory of epistemic supererogation. This will serve as a “proof of concept” – showing that a plausible theory of supererogation as a form of intrapersonal permissivism is at least possible. In all, this will show that there is good reason for intrapersonal permissivists to embrace epistemic supererogation.

## 2. *Intuitive Considerations*

Perhaps the most widely cited reason for thinking that rationality might be permissive is the intuitive existence of rational disagreement, even on the same evidence. Rosen (2001), for example, writes that when a jury is divided on a difficult case, or when paleontologists disagree about what killed the dinosaurs, it does not necessarily mean that someone has irrational beliefs. Kelly (2013) writes about a case where different agents disagree about how likely it is for a particular candidate to win the presidency before a close election. And Douven (2009) has argued that in some cases, scientists can rationally disagree about which theory best explains a body of evidence.

In these examples, we not only have intuitions about the existence of disagreement, but we clearly also have intuitions about what types of *evidence* engenders this type of disagreement. It is not a coincidence that these are cases where the evidence is extremely complex, multi-faceted, scarce, or fractured. Perhaps, then, certain bodies of evidence can have features that make them rationalize more than one doxastic response. When agents have this type of

permissive evidence, reasonable disagreement can happen. If it is the evidence that permits a range of rational responses, it is plausible to think that even a single agent with this sort of evidence can go either of two different ways, and end up rational. Thus, we end up with a natural way to start thinking about permissivism, and it is a way that is consistent with intrapersonal permissivism.

Though epistemic supererogation is rarely discussed, there is also something intuitive to be said for the idea. After all, many of the same considerations that motivate theorists to posit morally supererogatory actions also apply to certain beliefs. There exist beliefs that represent impressive feats of epistemic prowess, which seem to go beyond our epistemic requirements. Consider, for example, Einstein's theory of relativity, the proof of the existence of irrational numbers discovered by a member of Pythagoras' school, or the cases solved by fictional detective Sherlock Holmes. At least intuitively, these feats seem like epistemic analogues of the supererogatory acts of moral saints and heroes.

What is important to notice is that the cases where intrapersonal permissivism seems most intuitively plausible are also the cases where epistemic supererogation seems most intuitively plausible – cases where the evidence is complex, multi-faceted, scarce, or fractured. The particular nature of these bodies of evidence seems to be at least a partial explanation for why agents are not rationally required to respond to the evidence in the absolutely best way. This means that intrapersonal permissivism and epistemic supererogation mutually reinforce the plausibility of each other, in addition to whatever plausibility each view enjoys independently. So before we even get into the details of seeing why intrapersonal permissivists should embrace epistemic supererogation, we can see that a defender of either view already has some reason to



accept the other.

### *3. The Arbitrariness Objection*

To begin our examination of intrapersonally permissive theories, we will first consider a general problem for permissivism: the arbitrariness objection. If intrapersonal permissivism is true, then there is going to be some situation where an agent, given her evidence, can either believe  $P$  or believe  $\neg P$  and be rational in either belief. Here, however, we can begin to see an objection forming. For if both  $P$  and  $\neg P$  are rational to believe, then it seems that there is no reason to believe one over the other. Suppose that the agent knows all this about her epistemic situation. Then even from the agent's point of view, choosing one of these beliefs over the other seems awfully arbitrary.

To make this point vivid, imagine a rational agent who knows she is in a situation that is permissive between a belief in  $P$  and a belief in  $\neg P$ . She also has two pills – one which induces a belief in  $P$ , and one which induces a belief in  $\neg P$ . She could think hard about her evidence, weigh the different sides, and come to a belief based on that evidence, or she could randomly pick one of the pills and induce a belief in herself. Either way, she will end up with a rational belief. Thus, it seems like she has no reason to think that either method of belief formation is better than the other. But this is very unintuitive – clearly thinking about her evidence is rationally preferable to forming a belief at random. Indeed, after the agent forms a belief, she might as well take the pill for the contradictory belief. Since she knows this will also result in a rational belief, there seems to be no reason to think that there is anything wrong with changing her mind at will, rationally speaking. This situation seems clearly absurd. Yet it seems that

permissivists are committed to its possibility.<sup>12</sup>

One way to understand the worry behind this objection is to notice that there must be some connection between rationality and truth. Generally, when presented with a body of evidence we are warranted in believing that forming a rational belief given our situation is a good way to get to the truth. When an agent learns that her evidence is permissive, however, it seems that she also learns that the connection between truth and rationality is severed. To see this, suppose the agent considers a belief based on her evidence. Suppose she also knows that the belief is rational. With only this knowledge, it does not seem like she has good reason to think that the belief is likely to be true. After all, two incompatible beliefs are both rational, and only one of them can be true. So forming a belief in accordance with the evidence seems about as good as flipping a coin.<sup>13</sup> Thus, it seems that when she learns her evidence is permissive she also learns that her evidence won't do her much good. So an agent who realizes she has permissive evidence in regards to some proposition must suspend judgment on that proposition.<sup>14</sup>

Thus, if permissive situations are supposed to be situations where two different doxastic states are equally rational, then when agents realize this, there is going to be another doxastic state that is better than both – suspension of judgment. And if suspension of judgment is better than the “tied” doxastic states, then that is the rational doxastic state for the agent to form.

---

<sup>12</sup> Many variations of this example were first suggested by White in his (2005) and then his (2013). Other versions of this worry are presented by Christensen (2007) and Feldman (2007).

<sup>13</sup> Maybe the agent has some reason other than the belief's rationality to think that it is likely to be true. But it is incumbent upon the defender of permissivism to say what this other reason is. Later, we will see one attempt at meeting the challenge.

<sup>14</sup> See Feldman (2000, pp. 680), for a reason to think that in such cases, the uniquely rational doxastic response is suspension judgment. See Christensen (2007, fn. 8), for a discussion of the view that in similar cases, the uniquely rational doxastic response is representable by a range of credences between 0 and 1. I would like to thank an anonymous referee for pointing out these connections.

Which means that the situation is actually not permissive at all.

The arbitrariness objection is a roadblock that all intrapersonal permissivists must deal with. In general, there are two ways that permissivists have responded to it. In what follows, we will examine both strategies, and see why they both need to be complemented with a conception of epistemic supererogation.

#### *4. There Are No Cases of Known Permissivism*

##### *4.1. – The View*

Faced with the arbitrariness objection, one option for permissivists is to concede the basic line of reasoning. It *would* be absurd if an agent realized that she was in a permissive situation. So maybe such a realization destroys the permissivism. But this by itself does not rule out the existence of permissive cases – only cases of known permissiveness. The intrapersonal permissivist, then, can say that permissivism is true, but only in cases where the agent does not know she is in a permissive case.<sup>15</sup>

This option involves accepting the thesis that Cohen (2013) has called “doxastic uniqueness,” which says that an agent cannot rationally believe that there are two (or more) rational doxastic attitudes to take toward some proposition P, given total evidence E, while holding either doxastic attitude and having total evidence E.<sup>16</sup> There are two (compatible) ways to flesh out the details of doxastic uniqueness. One way is to think of the two doxastic states as “tied” or equally rational. The idea is that in some cases, an agent is permitted to form either of

---

<sup>15</sup> As we will see, multiple responses to the White paper can be considered versions of this view. White himself considers it briefly in his (2005). Brueckner and Bundy (2012) also explicitly discusses this general strategy of responding to White.

<sup>16</sup> Cohen (2013, p. 101)

two tied doxastic states, as long as she doesn't know they are tied.

The other model of permissivism involves thinking that permissive cases can happen even if one of the two permitted doxastic states is better supported than the other. This happens when, for whatever reason, the agent does not know that the rationally better belief actually is better. If this is rationally permissible, then the agent is permitted to form either the better belief or the worse belief. But it also seems that if the agent finds out the details of her situation – namely that some particular belief is better than some other particular belief – she must form the better belief. So again, when the agent realizes that her situation is permissive, her situation is no longer permissive.

There are different ways that one might fill in the details when attempting to deploy this general strategy. For illustrative purposes, let us quickly examine three of these attempts. I will illustrate the proposals with examples where the two permitted beliefs are tied, although they can also produce examples where one belief is better than the other.

The first possibility is suggested by Douven (2009).<sup>17</sup> Douven's proposal relies on the natural thought that, in many cases, some body of evidence supports a belief in some proposition because that proposition is a good explanation of the evidence. So for agents to form beliefs rationally, they have to think about how to explain the evidence they have. Some explanations, however, are extremely difficult to come up with – perhaps requiring nothing short of a brilliant flash of insight. But, Douven suggests, rationality may not *require* agents to have flash of insights. If this is right, then we have the possibility of permissive situations.

---

<sup>17</sup> Douven (2009, pp. 351-2), which features a permissive situation that arises due to some belief being supported by a brilliant flash of insight. Douven's situation is one where one belief is actually better supported than the other, but as already mentioned, I have adapted it to cases of "ties" for expositional purposes.

Imagine, for example, that a scientist has a large and complex body of evidence to consider. She attempts to come up with a good explanation for her evidence. As it turns out, there are two different and incompatible scientific hypotheses which explain the evidence equally well. Furthermore, both explain the evidence extremely well. The first hypothesis implies  $P$ , while the other implies  $\neg P$ . Both explanations, however, are extremely difficult to come up with – so difficult that rationality does not require the scientist to come up with either. In this situation, if the scientist comes up with the first explanation she is rational to believe  $P$ , but if she comes up with the second one she is rational to believe  $\neg P$ . So this situation is permissive. However, if the scientist comes up with both explanations, she will realize that her evidence supports both beliefs equally, and she can no longer either believe  $P$  or believe  $\neg P$ .

Another possibility, due to Rosa (2012), relies on the thought that sometimes bodies of evidence can be inconsistent.<sup>18</sup> According to Rosa, agents can have incompatible beliefs that they do not notice. In some of these situations, an agent can be rational in holding the inconsistent beliefs, and use them as evidence to form further beliefs. The idea is that these beliefs can constitute inconsistent bodies of evidence which might support both a belief in  $P$  and a belief in  $\neg P$ . For example, suppose that some agent has total evidence consisting of:

- (1)  $P$  and  $Q$
- (2) if  $P$  then  $R$
- (3)  $\neg Q$  or  $\neg R$

That is, the agent rationally has all three of these beliefs without realizing that they are inconsistent. With this evidence, she might reason to “ $R$ ” from (1) and (2). Or she might reason to “ $\neg R$ ” from (1) and (3). We might think that depending on which reasoning process she goes

---

<sup>18</sup> Rosa (2012, pp. 573-4)

through, the agent is permitted to either believe  $R$  or believe  $\neg R$ . Of course, if she goes through both reasoning processes, then she is not permitted to form either belief. Indeed, she needs to rethink her evidence.

Finally, Podgorski (2016) has suggested that which beliefs are rational for an agent might depend on how much of her evidence she takes into account.<sup>19</sup> And we might think that sometimes agents are not required to consider all of their evidence all the time. Suppose, for example, that an agent, receives two tiny bits of evidence regarding  $P$ . The evidence is small enough that it is overwhelmingly likely to make no difference as to whether  $P$ . Maybe in this case, the agent can neglect to think about these bits of evidence without being irrational. Of course, the agent *can* consider the pieces of evidence. Suppose that it just happens to be one of those cases where the small bits of evidence *do* make a difference as to whether  $P$ . In fact, if the agent just considers the first of the two bits of evidence, then she is rationally permitted to believe  $P$ . Had she considered just the second bit, she would be rationally permitted to believe  $\neg P$ . So this is a case where the agent could go either way, depending on which tiny bits of evidence she considers. Of course, if the agent knew all the facts about her situation – if, for example, she knew what would happen if she considered each bit of evidence – then she clearly is not permitted to form either belief.

#### 4.2. *The Problem*

Just looking at the different ways of fleshing out the view, we begin to see where the apparent problem lies. Each one of these permissive situations only happen when the agent

---

<sup>19</sup> See Podgorski (2016) §10 and §12. Podgorski's examples are most naturally interpreted as cases where one belief is actually epistemically superior to the other. Again, for ease of exposition, I have given a case where the two beliefs are "tied."

seems to exhibit some epistemic failing. The situations arise because the agent misses out on some explanations, reasoning processes, or considerations of evidence. Can any of this actually be rationally permissible?

One might think not. After all, evidential support relations seem knowable *a priori*. They are not the type of things that one typically gets empirical evidence for. Even if one did get such evidence, without knowing what the evidence supports (that is, without knowing the evidential support relations), the evidence would be no use. So at some point, agents have to be able to learn about these relations *a priori*. Thus, an agent with evidence that would permit believing either P or believing  $\neg$ P should be able to tell that this is the case. But this just means that she is able to tell that her evidence is permissive *a priori*. Thus, the agent with permissive evidence without realizing it is permissive is simply not accessing all the evidential support relations available to her. One might object that this in itself is less than rational. If that is right, then *any* rational agent who knows all the support relations must suspend judgment. So no agent can actually be rational in either believing P or believing  $\neg$ P, no matter what her evidence is like. In short, uniqueness turns out to be true.

More generally, the apparent problem for this view is that for an agent to be in a position where she would be rational to believe P and also rational to believe  $\neg$ P, the agent must be in a certain state of ignorance. This type of ignorance, however, is not due to some lack of empirical evidence on the agent's part. Clearly, agents can figure out what their evidence supports by reflection – so this knowledge is available *a priori*. The agent could have conquered her ignorance by better thinking alone. The agent thus seems epistemically culpable for her ignorance, or so one might argue. At the very least, it seems that an ideally rational agent would

not be so ignorant. Since it is intuitive to think that what an ideally rational agent would do is what we would be rational in doing, one might think that this strategy cannot really get us a permissive theory.

## 5. Epistemic Standards and Accuracy

### 5.1. The View

There is a different way to respond to the arbitrariness problem. This proposal employs the concept of “epistemic standards,” which we can think of as functions from bodies of evidence to doxastic states. Different standards encode different ways to respond to bodies of evidence. The general idea is that there is more than one standard that is rationally permitted. Each epistemic standard only advises a single attitude toward some proposition given a body of evidence. Different standards, however, might disagree about what this attitude is for some bodies of evidence. So there can be bodies of evidence where one rationally permitted epistemic standard will advise believing  $P$ , while another standard will advise believing  $\neg P$ .

This type of view can avoid the arbitrariness problem because it allows truth and rationality to come apart, from the point of view of any individual agent. This is because agents identify with the particular epistemic standard that they rationally believe to be the most reliable standard – the one most likely to get true beliefs.<sup>20</sup> From an agent’s own point of view, other

---

<sup>20</sup> Schoenfield (2014, p. 7) adopts this view of what it means to have a standard. Elga (ms, fn. 3) espouses a similar view. Part of the idea is that an agent’s epistemic standards, if they are rational, must be “immodest.” An agent’s epistemic standard is immodest just in case it advises beliefs that maximize *expected accuracy* from the agent’s own point of view, compared to the



agents using different standards may be rational, and the beliefs they end up forming in accordance with those standards are rational *for them*, but they are less likely to be true than the agent's own beliefs.<sup>21</sup> Thus, a rational agent can know that some situations are permissive. These are just the situations where her own standard disagrees with some other permissible standard, given her evidence. But choosing the belief her own standard advises isn't arbitrary for that agent, since she also thinks it is more likely to be true than the alternative.

This way of thinking, however, does not seem to result in intrapersonal permissivism. Since each epistemic standard only outputs one attitude toward P, intrapersonal permissivism is only true if agents are permitted to use more than one standard. It does not seem, however, that agents are ever permitted to do this. To see this, consider what agents believe about the reliability of their own standards. If agents rationally believe that their favorite standard is the most reliable one, then they cannot be rational in switching to what they think of as a less reliable standard. If they believe that their standard is less reliable than some other standard, then they were not rational in using that standard in the first place. Finally, if agents believe that all the standards are similarly reliable, we seem to run into the arbitrariness objection again. This is, at least, a natural line of thought – and the reason this strategy has mostly been seen as a path toward interpersonal, rather than intrapersonal, permissivism.

There is perhaps one way of avoiding this line of thought. Titelbaum and Kopec (ms.)

---

beliefs advised by any rival standard. Expected accuracy is a measure of how close an agent can expect some set of beliefs to be to the truth. Though this notion, favored by Bayesians, is usually defined in terms of credences, it is natural enough to discuss an analogous notion for all-or-nothing beliefs. See Lewis (1971) and Moss (2011) for examples of how to understand immodesty in terms of credences and the notion of expected accuracy. See Horowitz (2013) for a further discussion on immodesty as it relates to epistemic standards.

<sup>21</sup> See Schoenfield (2014, pp. 8-9) for a good presentation of this type of view. Subjective Bayesians also famously embrace this type of view.

offer a response to the arbitrariness argument which denies that once we realize some body of evidence is permissive, we have no reason to think that forming a belief on its basis is likely to get us a true belief.<sup>22</sup> Their view allows us to both say that rational epistemic standards are equally reliable at outputting true beliefs and that all standards are more likely to output true beliefs than false beliefs, *including* in permissive cases, where different standards output opposite beliefs.

To illustrate the proposal, let us focus on a toy example. Suppose that there are one hundred rationally acceptable epistemic standards and one hundred bodies of permissive evidence. For each body of permissive evidence, ninety of the acceptable standards output the true belief, and ten of them output the false belief. But for any acceptable epistemic standard, it outputs the true belief on ninety different bodies of permissive evidence, and the false belief on ten bodies of permissive evidence. Thus, even an agent who knows everything about the situation, including the fact that her evidence is permissive, can use her favorite rational standard and still be fairly sure (in fact, 0.9 sure) that the belief she ends up with is true.<sup>23</sup> So an agent can choose any rationally acceptable standard she wants as her favorite. Whatever she chooses, the beliefs she forms with it will very likely be true – including the beliefs she forms based on evidence *she knows to be permissive*. Thus, it seems that we have a view allowing us to avoid the arbitrariness objection while also allowing that all rationally acceptable standards are equally reliable – giving us the intrapersonal permissivism we were looking for.

---

<sup>22</sup> The view discussed here is especially inspired by §4 of their paper. Titelbaum and Kopec are not explicit on whether they mean their view to be a version of intrapersonal permissivism or not. So we can take the present view to be a version of intrapersonal permissivism inspired by their idea, which seems (at least initially) to be plausible.

<sup>23</sup> Titelbaum and Kopec (ms., p. 21)

## 5.2. *The Problem*

Titelbaum and Kopec argue that thinking about permissivism in this way affords us a nice way to maintain Conciliationism in the face of peer disagreement.<sup>24</sup> If one of these agents ran into her equally rational friend and found out that they disagreed about one of these propositions, even though they had the same evidence, then she is no longer rational in maintaining her belief. This is true even if the disagreement is because the agents are using different standards (and everyone knows this). This is because *after* she learns of the disagreement, she has no more reason to think it is her favorite standard that is the one getting it right. It could just as easily be her friend's. For example, suppose some agent rationally believes  $P$ , and finds out that her friend rationally believes  $\neg P$ , based on the same evidence. She now knows one of them is in the minority with the wrong belief, but she has no idea which one. So she must give up her belief and suspend judgment. So her new situation, which includes the evidence of disagreement, is not permissive. Instead, she must conciliate.<sup>25</sup>

While Titelbaum and Kopec see this as an advantage of their view, it also reveals a deeper difficulty. After all, why does the agent actually have to meet her friend in order to know that she would believe  $\neg P$ ? If she just knew what standard her friend was using, she could have figured out what her friend would believe (given that her friend didn't make a mistake). And in general, the agent should be able to figure out what a person would believe, given some specific epistemic standard. Indeed, if she knew all one hundred rational epistemic standards, then she

---

<sup>24</sup> Conciliationism refers to a broad family of views, according to which, a person should, upon discovering disagreement of a certain sort (normally disagreement about some proposition with someone with the same evidence and of similar intellectual prowess) revise her opinion in the direction of her disagreeer. Titelbaum and Kopec consider the view's compatibility with Conciliationism to be a nice feature of their view.

<sup>25</sup> Titelbaum and Kopec (ms., pp. 24-6)

could see which belief is in the majority for each body of potentially permissive evidence. Surely there is nothing in principle stopping her from doing this. Presumably, coming up with standards, figuring out which ones are rational, and determining which belief they recommend in her specific situation, is not something that requires empirical evidence. It is an *a priori* matter. But since the majority of the equally reliable standards always advise the true belief, if the agent did this she would have the true belief every time. So this “meta-standard” of consulting all one hundred rational standards and going with the majority is much more reliable than any one standard, even from the agent’s own point of view. One might object that the agent could not be rational in sticking to her favorite standard when she knows that all this is the case. And of course, all agents with the same knowledge of their situations should also use the meta-standard for the exact same reason. But this means all agents will be rationally required to come to the same beliefs, even given permissive evidence. In other words, it is hard to see how any of these bodies of evidence are permissive at all.

Once again, the apparent problem is that the purported permissive situations are ones where the agent is in a certain state of ignorance – in this case, ignorance about which standards are rational and what beliefs they would recommend. And again, since this ignorance is avoidable *a priori*, it does not seem to be a rational state to be in. If all this is right, then one might worry that we do not end up with a permissive theory at all.

### 6. *A Diagnosis*

We have seen two different attempts to construct intrapersonal permissive theories fall prey to essentially the same problem. In both cases, the purported situations in which an agent is

permitted to either believe  $P$  or believe  $\neg P$  only occur when the agents are in a certain state of ignorance. Furthermore, they seem to be states of ignorance that are *a priori* preventable. Since a defender of uniqueness could reasonably argue that it always seems irrational to be in a state of *a priori* preventable ignorance, these do not seem like situations that tell against the uniqueness thesis.

More abstractly, a uniqueness defender might argue that for any body of evidence, the possible doxastic responses to that body of evidence can be ranked according to their rationality. And this is a ranking that is determinable *a priori* – it is simply a matter of thinking up possible doxastic responses, figuring out how much they are supported by some body of evidence, and comparing their level of support. So every time an agent gets some evidence, she can just consult the ranking on that evidence. Once she has the ranking, it seems that she is rationally required to form the doxastic state at the top. After all, rationality is a guide to the truth. So, for any proposition  $P$ , if believing  $P$  is in the highest ranked doxastic state, then chances are  $P$  is true. And forgoing beliefs that are more likely to be true for beliefs that are less likely to be true is clearly irrational.<sup>26</sup>

For intrapersonal permissivism to be plausible, an alternative model must be offered. One possibility – that sometimes options are “tied” at the top of the ranking – has been blocked by the arbitrariness objection. In trying to escape this objection, both strategies we have canvassed rely on cases where the agent does not go through the entire procedure of constructing the ranking and finding the top doxastic state. In the first type of theory, the agent is in a

---

<sup>26</sup> Indeed, this might even be impossible, since it will likely involve believing  $P$  while also believing  $\neg P$  is more likely to be true. This is dangerously close to both believing  $P$  and disbelieving  $P$ .

permissive situation when she either does not know that her belief is tied with another one, or she does not know that another belief is ranked higher than her own. Under the Titelbaum and Kopec style view, the agent simply does not proceed to construct the ideal ranking, instead choosing to use a single epistemic standard and construct a non-ideal (but still fairly accurate) ranking. Without some explanation of how these situations are rationally permissible, both theories run into trouble.

### *7. Embracing Supererogation*

At this point, it is hopefully becoming clear why embracing supererogation is a good way forward for epistemic permissivism. The belief which results from coming up with the complete ranking and choosing the top doxastic state is clearly the maximally rational doxastic state for an agent to form. If there cannot be ties at the top of such rankings, then there is only one maximally rational doxastic state for any body of evidence. As long as epistemic agents are *required* to form the maximally rational doxastic state, uniqueness will be true. The only way out, then, is to claim that sometimes agents are not required to do what is maximally rational. This means opening the way for epistemic supererogation.

Let us see more concretely how epistemic supererogation will help. Borrowing from the literature on moral supererogation, we can work with a somewhat bare definition of what it means for a belief to be supererogatory. Namely, a belief is epistemically supererogatory just in case it is (1) not rationally required, (2) rationally permissible, and (3) rationally better than some alternative belief that is rationally permissible.<sup>27</sup>

---

<sup>27</sup> Plausibly, criterion (2) is redundant given criterion (3), since any belief that is rationally better than a permissible belief is itself permissible.

There are two ways we can employ epistemic supererogation in order to get us a theory of epistemic permissivism. We might say that, even after agents come up with the complete ranking of possible doxastic responses, agents are permitted to pick one that is not at the top. This means that there are some beliefs which are less than maximally rational, yet are permitted. So the maximally rational belief is (1) not required, (2) rationally permissible, and (3) rationally better than some alternative permitted belief (namely, the permitted non-maximally rational beliefs). However, for reasons already mentioned, I think it is implausible that agents can rationally form a belief that they know to be less rational (and therefore less likely to be true) than some alternative belief.

A different approach, which is more in line with the two strategies examined in this paper, is to claim that agents can sometimes permissibly fail to figure out the complete ranking. That is, agents are sometimes allowed to be in states of partial ignorance about the ranking. Figuring out the complete ranking would be rationally better, of course. But on the present view, it would be supererogatory.

However we ultimately develop it, many of the theories we have examined become much more plausible when embedded within a larger theory of epistemic supererogation. We might think, for example, that not being able come up with certain complex explanations is rationally permissible, but being able to come up with such hypotheses is much better. Thus, the beliefs resulting from the rationally better explanations are supererogatory. Or perhaps seeing that one's evidence is inconsistent is not always required, given that the inconsistency is hard enough to see. But seeing the inconsistency is better – and hence, results in supererogatory beliefs. Alternatively, agents might be permitted to not always consider all of their evidence, although

the epistemic saints who do are doing better, epistemically speaking. And finally, maybe ordinary epistemic agents only use one reliable epistemic standard, and this is okay. But if they considered all rational standards, they would be doing much better than okay – they would be epistemic heroes.

This is not to say that, on the final analysis, all of these views can be developed into a successful theory of supererogation. But every theory we have so far considered gives us a way into such a theory, if we are only willing to head in that direction.

### *8. Proof of Concept*

We have seen that a theory of epistemic supererogation can solve the problems that intrapersonal permissivism has with the arbitrariness objection. This seems like reason for intrapersonal permissivists to embrace epistemic supererogation. This is the main conclusion of the paper. However, a skeptical permissivist may still have doubts. After all, this advice is completely useless if no plausible theory of epistemic supererogation can be constructed. Moreover, even if we can develop a plausible theory of epistemic supererogation, not just any theory will do – we need a theory which allows for rationally permissible states of *a priori* preventable ignorance. Finally, simply asserting that such doxastic states are permissible would do little to answer the original objection. Thus, a theory of supererogation needs to explain, in a plausible way, why these states of ignorance get a rational pass. We might worry that all this is too tall of an order for a theory of epistemic supererogation to actually reach.

To alleviate this worry, I will briefly present the theory of epistemic supererogation I develop in detail in my (Li 2017). Armed with this theory, we will be able to tell a plausible



story about why sometimes agents are rationally permitted to be in states of *a priori* avoidable ignorance. Though the main conclusion of this paper will not depend on the details of any specific theory, this section will serve as evidence that it is at least possible to construct a theory can do the work interpersonal permissivists require of epistemic supererogation.

To begin constructing a theory of epistemic supererogation, we can look toward extant theories of moral supererogation for guidance. Specifically, some theorists contend that moral supererogation happens because actions can be judged with regards to two different types of moral virtues or values: the value of justice and the value of beneficence.<sup>28</sup> Justice can require certain actions of agents, while beneficence can only justify actions. Thus, certain beneficent actions are morally supererogatory because the value of beneficence cannot generate moral requirements. This rough story can be transferred to the epistemic realm if we can find two epistemic virtues, only one of which can generate rational requirements.

To find such virtues, we can look back toward the intuitive cases of epistemic supererogation discussed in section 2. Consider the case of Einstein's discovery of the theory of relativity. What is most impressive about Einstein's achievement was his ability to come up with such a radically different theory – with its fundamental revisions of our concepts of space and time – that fit all the data. His resulting belief was so surprising that epistemic agents can be forgiven for overlooking its possibility – that is why it seemed supererogatory in the first place. Something analogous is true of our other putative examples of epistemic supererogation. Sherlock Holmes, for example, was such a great detective precisely because he was able to think up the incredibly unusual (but ultimately correct) explanations for his evidence.

---

<sup>28</sup> Theories with this idea at their core have been proposed by Zimmerman (1993) and Dreier (2004), although here I am not ascribing the view to anybody in particular.

Notice, however, that Einstein was not necessarily better at evaluating how well a given hypothesis is supported by the evidence. Indeed, even an undergraduate physics student can understand the theory of relativity and why it makes sense of the evidence. For these students, there is no intuition that such a belief would be supererogatory. If the physics student fails to see how the theory of relativity is supported by the data even after being suitably taught, then she is being irrational. Intuitively, this is because the hard part – coming up with hypothesis – has already been done for them.

From these observations, we can distinguish two different epistemic virtues. One is the more everyday virtue of seeing the support relationships between evidence and hypotheses. This is more of a housekeeping virtue, requiring something like analysis and critical reasoning. The other is the virtue of coming up with the hypotheses themselves, requiring more creativity and imagination.<sup>29</sup> The rough proposal, then, is that doxastic states can be evaluated in regards to whether they exhibit the creative virtue. Though beliefs that exhibit this virtue are epistemically better for it, the creativity is never required. Only housekeeping considerations can require doxastic states.

Thus, the creativity involved in coming up with the theory of relativity explains why such a belief was not required, and in fact supererogatory, for Einstein. In order for this explanation to make sense, however, we need to refine our understanding of what it means for a doxastic state to be required. This is because, in the first instance, it was the process of belief formation that was an exercise of the creative virtue – not the belief itself. After all, it was this process that involved actually coming up with the radically different hypothesis. Therefore, it was the

---

<sup>29</sup> For a different discussion of this same distinction, see Nozick (1993), starting on p. 172.

process of belief formation that was not required, and in fact, supererogatory. If the same belief was formed by another process – say just being told about the theory by a physics teacher – the belief would not be supererogatory.

The idea of epistemically evaluating something other than a doxastic state is not foreign to epistemology. Indeed, one of the few extant discussions of epistemic supererogation relies heavily on this idea. In his (2014), Hedberg discusses cases of non-required, and therefore epistemically supererogatory, *acts*. His examples include the act of double checking one's evidence and the act of gathering additional evidence. The theory presented in this section is sympathetic to Hedberg's general strategy of finding epistemic supererogation in processes that are not epistemically required. However, readers who are skeptical that acts such as acquiring evidence can be epistemically evaluated may feel more comfortable with the thought that the process of coming up with a hypothesis is epistemically evaluable. This process is, at least, a cognitive process that does not change one's evidential state.

In any case, Hedberg's treatment of epistemic supererogation does not extend to actual doxastic states. For our theory of epistemic supererogation to be a version of intrapersonal permissivism, however, we must take this extra step. We need it to be the case that Einstein's actual *belief* in the theory of relativity was not epistemically required. To get this result, we can simply claim that whether a belief is epistemically required is parasitic on whether the belief forming process that created it was epistemically required. So we can say that Einstein's belief was not required in this sense: it was the result of a belief forming process that itself was not required.

This sense of a non-required doxastic state is derivative – we cannot tell, for example,

whether a belief was required unless we know how the belief was formed. But this type of derivative epistemic property is also not foreign to epistemology. The standard way to understand a doxastically justified belief – a belief that is propositionally justified and also formed in the “right way” – is one such example. This is another case where the epistemic status of a belief forming process (being formed in the “right way”) confers a certain epistemic status (doxastic justification) upon the belief that the process produces. Notice, for example, we also cannot tell whether a belief is doxastically justified unless we know how the belief was formed. If there is no problem understanding doxastic justification as a property of a belief, there should also be no issue with this conception of a required belief.

So Einstein’s belief was not rationally required. But it was also clearly rationally better than some alternative rationally permitted belief (namely, suspension of judgment). So Einstein’s belief is supererogatory. Einstein’s contemporaries, however, were perfectly rational in not believing in relativity. This is because coming up with the theory of relativity was itself not required. More generally, this is why it is sometimes rationally permissible to be in a state of *a priori* preventable ignorance. If the only way to avoid the state of ignorance is to engage in some non-required exercise of the creative virtue, then such ignorance is rationally permissible. In essence, the theory claims that sometimes coming up with a complete list of possible doxastic responses to a body of evidence is not required, since some of those possible responses would take supererogatory acts of creativity to even entertain. Thus, in these cases, rational agents are not required to construct the entire ranking of doxastic responses in terms of their plausibility.

Agents who nevertheless construct the entire ranking may end up with a supererogatory belief.<sup>30</sup>

Since this theory can explain why some cases of *a priori* preventable ignorance is rationally permissible, we have a form of intrapersonal permissivism that avoids the arbitrariness objection. This theory also makes sense of our intuitions in many paradigmatic cases of seeming epistemic supererogation. Thus, there is good reason to think that the theory will be very appealing to an intrapersonal permissivist. Of course, there may be other theories that can perform the same task – it is beyond the scope of this paper to consider them all. We have seen, however, that a theory which seems to meet all the permissivist’s desiderata is possible.

### 9. Conclusion

We have seen that not only is epistemic supererogation a form of intrapersonal permissivism, it is perhaps our best hope for developing a plausible theory of this type. At least two promising strategies of developing intrapersonal permissivism turn out to suffer a common defect – the purportedly permissive situations they posit all require agents to be in a state of *a priori* preventable ignorance. On a natural way of thinking about epistemic rationality, this is not a rational state to be in. The best way to patch up this defect is by creating conceptual space for a less than maximally rational, but still rationally permissible, doxastic state. In short, it requires the existence of epistemic supererogation. So our best hope for a theory of intrapersonal permissivism rests in a theory of epistemic supererogation. We have also seen that a plausible theory of supererogation is at least possible. Combine this with the independent plausibility of

---

<sup>30</sup> As this has just been a summary of the view, there are many issues left unresolved. It is unclear, for example, just what it means to “come up” with a new hypothesis, and when doing so is not epistemically required. These issues are discussed in more detail in Li (2017).

each view considered in isolation, and with the mutual support the views lend each other when considered together, we end up with a package that has much to recommend it.

*Han Li, Kansas State University*

Acknowledgements: For comments on drafts of this paper, as well as discussions about these and related issues, I would like to thank an anonymous reviewer for *American Philosophical Quarterly*, Zachary Barnett, David Christensen, James Fritz, Christopher Meacham, Bradford Saad, Joshua Schechter, Paul Silva, participants of the 2014, 2015, and 2016 Dissertation Workshop at Brown University, and the audience at my presentation during the meeting of the Eastern Division of the American Philosophical Association in 2017.

#### References

- Ballantyne, Nathan, and E. J. Coffman. 2011. "Uniqueness, Evidence, and Rationality." *Philosophers' Imprint* 11 (18).
- Berker, Selim. 2013. "Epistemic Teleology and the Separateness of Propositions." *Philosophical Review* 122 (3):337–393.
- Brueckner, Anthony, and Alex Bundy. 2012. "On 'Epistemic Permissiveness.'" *Synthese* 188 (2):165–177.
- Christensen, David. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review* 116 (2):187–217.
- . 2016. "Conciliation, Uniqueness and Rational Toxicity." *Noûs* 50 (3):584–603.
- Cohen, Stewart. 2013. "A Defense of the (Almost) Equal Weight View." In *The Epistemology of Disagreement*, edited by David Christensen and Jennifer Lackey, 98–119. Oxford University Press.  
<http://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199698370.001.0001/acprof-9780199698370-chapter-6>.
- Douven, Igor. 2009. "Uniqueness Revisited." *American Philosophical Quarterly* 46 (4):347–361.
- Dreier, James. 2004. "Why Ethical Satisficing Makes Sense and Rational Satisficing Doesn't." In *Satisficing and Maximizing*, edited by Michael Byron. Cambridge University Press.
- Elga, Adam. n.d. "Lucky to Be Rational."
- Feldman, Richard. 2000. "The Ethics of Belief." *Philosophy and Phenomenological Research* 60 (3):667–95.
- . 2007. "Reasonable Religious Disagreements." In *Philosophers Without Gods: Meditations on Atheism and the Secular*, edited by Louise Antony, 194–214. OUP.
- Hedberg, Trevor. 2014. "Epistemic Supererogation and Its Implications." *Synthese* 191 (15):3621–3637.
- Horowitz, Sophie. 2014. "Immoderately Rational." *Philosophical Studies* 167 (1):41–56.
- Kelly, Thomas. 2013. "Evidence Can Be Permissive." In *Contemporary Debates in Epistemology*, edited by Matthias Steup and John Turri, 298. Blackwell.
- Lewis, David. 1971. "Immodest Inductive Methods." *Philosophy of Science* 38 (1):54–63.

- Li, Han. 2017. "A Theory of Epistemic Supererogation." *Erkenntnis*, March, 1–19.  
<https://doi.org/10.1007/s10670-017-9893-3>.
- Moss, Sarah. 2011. "Scoring Rules and Epistemic Compromise." *Mind* 120 (480):1053–1069.
- Nozick, Robert. 1993. *The Nature of Rationality*. Princeton University Press.
- Podgorski, Abelard. 2016. "Dynamic Permissivism." *Philosophical Studies* 173 (7):1923–1939.
- Rosa, Luis. 2012. "Justification and the Uniqueness Thesis." *Logos and Episteme*, no. 4:571–577.
- Rosen, Gideon. 2001. "Nominalism, Naturalism, Epistemic Relativism." *Noûs* 35 (s15):69–91.
- Schoenfield, Miriam. 2013. "Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief." *Noûs* 47 (1):193–218.
- Titelbaum, Michael G., and Matthew Kopec. n.d. "Plausible Permissivism."
- Urmson, J. O. 1958. "Saints and Heroes." In *Essays in Moral Philosophy*, edited by A. I. Melden. University of Washington Press.
- White, Roger. 2005. "Epistemic Permissiveness." *Philosophical Perspectives* 19 (1):445–459.
- . 2013. "Evidence Cannot Be Permissive." In *Contemporary Debates in Epistemology*, edited by Matthias Steup and John Turri, 312. Blackwell.
- Zimmerman, Michael J. 1993. "Supererogation and Doing the Best One Can." *American Philosophical Quarterly* 30 (4):373–380.