How Reasons Balance

Ethan Kemp

1 Introduction

What determines what I should do? What determines what I should believe? Here are two platitudes. (1) I should do what I have most reason to do; (2) I should believe what I have most reason to believe. Both of these statements appear to reflect the same underlying truth: reasons settle which choices are correct. Accordingly, discussions in recent normative philosophy have increasingly focused on the need for *unified* theories of both practical and epistemic reasons. Additionally, much has also been written on the *commensurability* of these two kinds of reasons – that is, whether practical reasons can be rationally weighed against epistemic ones.[1]

But what should I do when my reasons for acting in a certain way, or believing a particular proposition, are tied with competing reasons? What ought I to do when my reasons are *balanced*? A common answer to this question is that it depends on whether the reasons are practical or epistemic. Simply put, when agents have exactly equal practical reason for φ-ing as they do for ψ-ing (when practical reasons are balanced), we think they are rationally permitted to φ, and rationally permitted to ψ. In other words, practical reasons exhibit *permissive balancing*. By comparison, when epistemic reasons balance (when agents have exactly equal reason to believe p as they do to believe $\neg p$), we think agents are rationally prohibited from believing p, and rationally prohibited from believing $\neg p$.[2] Epistemic reasons thus appear to exhibit *prohibitive balancing*.[3] The common response, then, is that we can't give a general answer to the question of what I ought to do when my reasons are balanced, since the answer depends on which kind of reason is in play.

This is often assumed to mark a crucial difference between practical and epistemic reasons.[4] For example, Harman (2004), Feldman (2006), Cohen (2016), and Dancy (2018) suggest this asymmetry poses a serious challenge for unified theories of the practical and epistemic domains. Others, like Berker (2018), exploit this difference to argue that practical

[1] Though the literatures for these two topics focus on different questions, they are closely intertwined. For discussion on the unification of practical and epistemic reason, see Raz (1999), Scanlon (1998, 2014), Dancy (2004, 2018), Skorupksi (2010), Parfit (2011), Horty (2012), Schroeder (2021). For literature on the commensurability of practical and epistemic reasons: Kelly (2002), Hieronymi (2005), Reisner (2008), Rinard (2017, 2019), Berker (2018).

[2] In this paper I will use "believing $\neg p$", and "disbelieving p", interchangeably.

[3] The "permissive" and "prohibitive" balancing terminology is borrowed from Berker (2018).

[4] Feldman (2006, 229)

and epistemic reasons are like "oil and water" – that they are entirely incommensurable.[5] Certainly, if practical and epistemic reasons operate in essentially different ways, we should be skeptical of the possibility for unification and commensurability.

My target is this assumption; the assumption that practical reasons and epistemic reasons balance in essentially different ways – and that this explains our observations about what it is rational for agents to do or believe in practical and epistemic cases. I call this the *Balancing Assumption*. I believe this assumption is not warranted, and so I intend to replace it with a better explanation of balancing behavior. I argue that cases of so-called prohibitive balancing are, instead, cases where a single unique response is most-favored by one's combined reasons. For example, in epistemic cases where one's reasons for believing p are balanced with their reasons for believing $\neg p$, their reasons *most-favor* being agnostic about p.[6] If I am right, then cases of prohibitive balancing are, in fact, not cases of balancing at all. I think such an account has intuitive force: what we ought to do or believe in a given situation is just what we have most reason for. Since we ought to be agnostic when our reasons for p and $\neg p$ are balanced, we must have most reason for being agnostic.

My proposal, then, is that cases of permissive balancing are cases where our reasons fail to pick out a unique response, and cases of (merely apparent) prohibitive balancing are cases where our reasons do pick out a unique response. This, of course, stands in need of further explanation. It also generates a puzzle that we will need to resolve. The puzzle is this: how can it be that when we have reason for p, reason for $\neg p$, and *no other reasons*, that we really have *most reason* for being agnostic?

In §2, I lay out and discuss canonical cases of apparent permissive and prohibitive balancing, and offer my preferred explanation. In §3, I outline the puzzle that arises in virtue of explaining cases in this way, and reject one potential way of resolving it. I offer my own solution in §4. In order to defend this solution, I develop an account of how reasons favor actions/beliefs, which I call the *Structural Account*.

The result will be valuable to both pragmatists (roughly, those who are friendly to unification and commensurability), and to evidentialists (those who are skeptical of unification and commensurability).[7] Though I don't expect to settle such long-standing debates here, I hope to shift the focus of such debates away from the Balancing Assumption, which attempts to explain balancing phenomena in terms of essential properties of reasons in-themselves.

[5] Berker (2018, 434)

[6] In this paper I assume a standard picture where there are three doxastic attitudes rational agents can take towards a proposition: belief, disbelief, and agnosticism. I follow Friedman (2013) in thinking that *being agnostic* is its own positive attitude, distinct from merely lacking belief in p and in $\neg p$. On this view, to be agnostic whether p is to have, roughly, a currently-neutral doxastic attitude towards p. Importantly, McGrath (2021) has argued that nearby concepts like *suspending*, *withholding* or *refraining* from judgement, connote ways of *acting* to delay judgement-formation, which differs from simply having a doxastic attitude of neutrality towards a proposition. This distinction will be useful later in this paper.

[7] "Evidentialist" is the standard name for those who think there cannot be practical reasons for belief, and "pragmatist" is the standard name for those who think there are (or can be). I use these labels in this paper to refer more broadly to the loose camps of philosophers who are skeptical of unification and commensurability, and friendly to unification and commensurability, respectively.

In §5, I examine a potential objection: if epistemic reasons do not balance prohibitively, then this account seems to entail the controversial thesis of *permissivism*.[8] There are two ways of responding. The first is to make permissivism merely vacuously true by arguing that epistemic reasons never find themselves in a state of balance. This way of responding suggests more work to be done, and I'll discuss one worry that comes along with it in §6. The second way of responding (which I am partial to), is to accept that there are genuinely-permissive epistemic cases. In the end, I find both ways of responding satisfactory, and the reader's preferred route will depend on their prior theoretical commitments. Regardless, I will have shown that the Balancing Assumption is worth abandoning.

2 Cases and Challenges

To better grasp why the Balancing Assumption strikes many philosophers as plausible, let's first consider examples that seem to demonstrate practical reasons balancing permissively, and epistemic reasons balancing prohibitively. The following two cases I adapt from Brunero (2022):

> **Party or Library** Charlie is choosing whether to go to a party to catch up with old friends, or to instead head to the library to study for his exam tomorrow morning. He knows that he can't do both. Suppose he has good reasons to go to the party, and good reasons to go to the library, and neither set of reasons outweighs the other. Assuming that there are no other relevant reasons in play, it's permissible for Charlie to decide to go to the party, and permissible for Charlie to decide to go to the library. This is a case of permissive balancing.

> **Will He Attend** Dani is deliberating about whether Charlie will show up at the party. She has some good reasons for thinking that Charlie will show up, but also some good reasons for thinking he won't. The reasons for thinking Charlie will show up do not outweigh, nor are outweighed by, the reasons supporting Charlie not showing up. In light of this, it's not permissible for Dani to believe that Charlie will show up, and not permissible for Dani to believe that Charlie won't show up. This is a case of prohibitive balancing.

In the first case, Charlie is deliberating about whether to go to the party or to the library. His reasons are practical. He has exactly equal, or balanced, reasons for each choice. Given this balance, it is generally recognized that Charlie is rationally *permitted* to engage in either action. There is nothing irrational about his doing one or the other. In other words, his reasons exhibit permissive balancing. Conversely, Dani has reasons for believing that Charlie will show up, and exactly equally-weighed reasons for believing that he will not show

[8] By this I mean a conditional version of permissivism: *if epistemic reasons are ever in a state of balance, then more than one doxastic state will be rationally permissible*. This allows for the possibility of a vacously-true permissivism. I will discuss this in §5.

up. Dani's reasons are also balanced. However, Dani's reasons are epistemic. It is widely agreed that the rational response to balanced epistemic reasons like Dani's is to be agnostic on the matter. That is, we think that Dani is rationally *prohibited* from believing either of her options – she should instead remain agnostic about whether Charlie will or won't show up to the party. Her reasons, therefore, appear to exhibit prohibitive balancing.

Though Charlie and Dani are in the same situation with regard to the balance of their reasons, one of them is permitted to engage in either action, while the other is prohibited from believing either proposition. This difference is puzzling, but it is *prima facie* natural to assume that the relevant difference here has to do with Charlie's reasons being practical, and Dani's reasons being epistemic. One may be tempted (as many are) to accept this explanation as complete: since practical reasons balance permissively, and epistemic reasons balance prohibitively, this explains the difference in which choices are permissible for Charlie and Dani. To opt in to such an explanation is to essentialize the permissive and prohibitive natures of practical and epistemic reasons, respectively.[9]

As I suggested in the introduction, taking practical and epistemic reasons to have essential differences in balancing behavior inevitably creates difficulties for unification and commensurability. It's a small move from essential differences between practical and epistemic reasons, to the view that the practical and epistemic domains are disunified and incommensurable.

To illustrate this move, consider Schroeder's (2015) theory of what makes reasons sufficient. On his account, a set of reasons is sufficient just in case it is at least as weighty as any other competing set of reasons. Schroeder's intention is to provide a unified theory of *sufficiency* which applies to both practical and epistemic reasons. However, the Balancing Assumption challenges such an account. In cases where epistemic reasons are balanced, they meet Schroeder's condition for sufficiency – balanced epistemic reasons are *at least as weighty* as the competing reasons they are balanced against. But if it is true that epistemic reasons balance prohibitively, then we don't think an agent with balanced reasons for p and $\neg p$ has sufficient reason to believe p or sufficient reason to believe $\neg p$. As Cohen (2016) puts it, "it can be rational for one to act, even if one has only a weak reason, provided one does not have a stronger reason not to act. Belief is different. I am not rationally permitted

[9]One potential way of supporting this essentializing line of thought, which I want to set aside, is to appeal to the disanalogy between action and belief. The disanalogy lies in the fact that for belief, agents always have the option of being agnostic, which is not the case for action. The thought continues that this disanalogy between action and belief explains why practical reasons balance permissively, and epistemic reasons balance prohibitively. However, I do not believe this is sufficient to explain the behavior of practical and epistemic balancing. First: it's not clear how strong the suspension of judgement disanalogy actually is. Buridan's Ass, for example, seems to have three options: go to the hay bale on the left, go to the hay bale on the right, or do nothing. This seems analogous to doxastic options (believe, agnostic, disbelieve), and yet Buridan's Ass's practical reasons still give him rational permission to choose one or the other hay bales. Second, as noted by Berker (2018), practical reasons *for belief* (if they exist), would nevertheless balance permissively. If an agent, for example, was offered \$1,000 to believe p, and \$1,000 to believe $\neg p$, these practical reasons would be permissively balanced between the two options. Therefore, it does not seem that we can cash out the differences in practical and epistemic balancing simply in terms of the differences in action and belief.

to believe that a fair coin toss will land heads." So, Schroeder's account is an example of a *unifying* theory which is at risk of being false if the Balancing Assumption turns out to be true.

The Balancing Assumption also casts doubt on the commensurability of practical and epistemic reasons. To see why, imagine that you have a practical reason for φ, and an equally-weighed epistemic reason for ¬φ. If practical reasons always balance permissively, does this mean, in this scenario, that your practical reason gives you permission to φ and permission to ¬φ? Conversely, does your epistemic reason prohibit you from φ-ing and from ¬φ-ing? It seems in such a scenario that our practical and epistemic reasons conflict in a way that gives no rational guidance on how to respond. Berker (2018) develops these counterintuitive interactions further, arguing that any way we might try to weigh practical and epistemic reasons against one another yields seemingly irrational verdicts, leading him to conclude that practical and epistemic reasons are simply incommensurable.

I will return to each of these worries in §6. Consider now a curious third case:

> **Shy Friends** Enzo is enjoying a party and deliberating about whether to go sit next to his shy friend Ava, or his shy friend Bob. There is a seat directly next to Ava, and a seat directly next to Bob. Sitting in the seat next to Ava would preclude Enzo from talking to Bob, and vice versa. However, there is a third seat, equidistant from each, which would allow Enzo to talk to both Ava and Bob, although not quite as easily as if he was sitting right next to one or the other. Enzo has good reasons for sitting next to Ava, and equally good reasons for sitting next to Bob (Ava and Bob each look miserable without a conversation partner), and no other relevant reasons. In this case, it doesn't seem permissible for Enzo to sit directly next to Ava, nor does it seem permissible for Enzo to sit directly next to Bob. Instead, he must sit in the third seat, in between his two friends, so that he can talk to both. This is a case of prohibitive balancing.

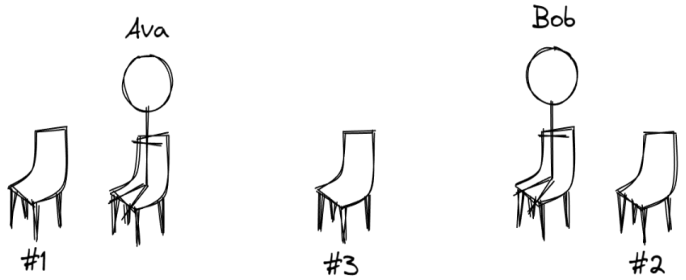

In this example, Enzo is deliberating about where to sit at the party. Enzo's reasons are practical. If Bob was not present, then Enzo would sit directly next to his friend Ava, and if Ava was not present, Enzo would sit directly next to Bob. This suggests that Enzo has a reason for sitting in chair #1, a reason for sitting in chair #2, and no other reasons. Considering his equally-weighed reasons, it seems that Enzo is *prohibited* from sitting directly

next to A, and *prohibited* from sitting directly next to B. If practical reasons always balance permissively, as the Balancing Assumption would have it, then Enzo would be permitted to sit directly next to Ava, and permitted to sit directly next to Bob. But it looks as if Enzo is not permitted to do either of these things. Enzo's reasons thus appear to exhibit prohibitive balancing, even though they are practical reasons.

If this is a genuine case of practical prohibitive balancing, then it immediately weakens support for the Balancing Assumption. Recall that a compelling reason for thinking there is an essential difference between the way practical and epistemic reasons balance is that practical reasons seem to *always* balance permissively (and epistemic reasons always prohibitively). If Enzo's practical reasons are balanced prohibitively, though, we cannot say it is the essential nature of practical reasons to always balance permissively.

One might think, however (as I do), that **Shy Friends** is not a genuine case of prohibitive balancing. On this line of thought, Enzo just has most reason – more than any other competing reasons – to sit in chair #3. Therefore, by describing Enzo's reasons as "balanced", we have misdescribed the case. Certainly, Enzo's reasons for sitting in chair #1 and for sitting in chair #2 are equal. But both of these options have less overall reason in their favor than sitting in chair #3. So, by saying Enzo's reasons are balanced, we fail to identify the option which is favored most.

I think Dani's situation is closely analogous to Enzo's. That is, when we say epistemic reasons balance prohibitively, we make the same failure of description that we do when we say Enzo's practical reasons are balanced prohibitively. In epistemic cases where we have equal reason for believing p as we do for $\neg p$, we also have *most* reason for being agnostic whether p. So-called cases of epistemic prohibitive balancing, then, are not really cases of balancing at all – they are cases where agnosticism is the uniquely-favored response.[10] In the following two sections, I will defend this idea.

3 A New Puzzle

Our initial description of **Will He Attend** was that Dani has reason for believing Charlie will attend the party, equal reason for believing Charlie will not attend the party, *and no other reasons*. In order to say – as I want to – that Dani actually has most reason for being agnostic, we must now resolve a puzzle: how can it be that Dani has most reason for being agnostic, if, when describing her set of reasons, it seems she has *no reasons at all for being agnostic*?

There is an existing account, developed by Schroeder (2021), which may provide a way of addressing this puzzle. The strategy? Deny that there are "no more reasons". Schroeder has argued that there are epistemic, but non-evidential, reasons for suspending judgement

[10]Notice that if we accept the Balancing Assumption, we have to accept a surprising implication: in the case of belief, there are times (when your reasons are balanced) that you *shouldn't do what you have most reason to do*. This follows directly from saying that epistemic reasons balance prohibitively. And this should strike us as implausible.

in belief-cases like Dani's. He writes: "what drives all of our puzzles is the assumption that the only reasons that matter for the rationality of belief are evidence for or against its content… What I want to suggest, is… there are always reasons against belief that are not evidence".[11] The idea is, when one's evidence for and against a proposition is tied, there are also additional reasons in favor of being agnostic, which ultimately makes it the most-favored attitude.

Let's call this the *Extra Reasons* strategy for resolving the puzzle. While there is much to be said in its favor, I think it also faces significant challenges. In service of developing my own novel strategy, I'll briefly sketch some of these challenges before moving on to develop my preferred way of resolving the puzzle.

First, is that Extra Reasons does not explain practical cases like **Shy Friends**. Enzo has a reason for sitting in chair #1, a reason for sitting in chair #2, and no other reasons. What explains the fact that Enzo has most reason for sitting in chair #3? Extra Reasons only offers a resolution for epistemic cases, but not for practical ones. Ideally, a strategy for resolving one will resolve both.

Even if we focus solely on epistemic cases, a second objection goes: it is not obvious that agents always have non-evidential epistemic reasons against belief. Some examples Schroeder offers are: (1) if forming the belief would be useless in forming further inferences, (2) the practical costs of belief, (3) the fact that more evidence is coming in the future. But notice that in **Will He Attend,** it is not obvious that Dani has any of these reasons.[12] A belief about whether Charlie will attend the party would clearly be useful for forming further inferences. Dani doesn't have any real stakes in her belief about whether Charlie will attend, so there's no practical cost. And we can just assume for Dani's case that there is no evidence coming in the future. Nevertheless, we think Dani rationally ought not believe Charlie will attend the party (and ought not believe he won't attend). Whatever the non-evidential epistemic reasons are, they would have to be reasons we *always* have – and it's not obvious that there are such reasons.

Third objection. Assume for sake of argument that there's an omnipresent reason against believing. What explains how this reason always "scales up" to outweigh reasons for belief, regardless of the strength of the balanced reasons for p and $\neg p$? Imagine that I have weak reason for believing p, and weak but equal reason for believing $\neg p$. On Schroeder's model, my extra reason for suspending judgement outweighs both of these reasons, and so I ought to be agnostic. However, imagine that I later obtain strong evidence for p, and equally strong evidence for $\neg p$. My non-evidential reason for suspending judgement has not increased in strength, and yet, on Schroeder's model, it somehow still outweighs my strong reasons for p and for $\neg p$. It's not clear what explains this.[13]

A fourth challenge is worth mentioning. McGrath (2021) has argued that reasons like future-incoming-evidence bear on whether one ought to suspend judgement whether p, but

[11] Schroeder (2021, 141)
[12] Brunero (2022, 92)
[13] See Berker (2018) and Brunero (2022) for discussion of this objection.

not on whether one should be agnostic towards p. This is an important difference. On McGrath's taxonomy, to suspend judgement is to put off deliberation, or "belief-forming judgement" until a later time. This appears to be a kind of action, rather than an attitude towards a proposition. Even if Schroeder is right, then, that we always have some extra reason to delay belief-formation, it's not obvious that incoming future evidence bears on whether you're currently doxastically neutral on p, in the strict sense of *being agnostic* with which we have been talking.

I have just offered four considerations against the Extra Reasons response to our puzzle about how agnosticism can be most-favored in cases like **Will He Attend**. Since I take it that Extra Reasons will not be convincing to those who are already partial to the view that practical and epistemic reasons balance in essentially different ways, I'll set aside this strategy in order to develop my own. In doing so, I have two aims. First, the strategy should be independently plausible. This will ward off *ad hoc*ness concerns. Second, it should maintain as much theoretical neutrality as possible between practical and epistemic reasons.

4 A New Model of Reasons' Favoring

To reject the Balancing Assumption, we must explain how cases of apparent prohibitive balancing are instead cases where agents have most reason in favor of one unique thing. The description of such cases – e.g. that agents have reason for p, reason for $\neg p$, and no other reasons – makes this especially challenging. In this section, I offer my own way of resolving the puzzle: balanced reasons, *in combination*, can most-favor some unique thing, despite not *individually* most-favoring that unique thing.

Imagine (for simplicity's sake) that Dani has just one reason, E_1, in favor of believing that Charlie will show up to the party, and just one equally strong reason, E_2, in favor of believing that Charlie will not show up to the party. On the account I want to embrace, E_1 most-favors Dani's believing that p. If Dani's only relevant reason was E_1, then we would say she ought to believe that p. The vice versa is true of E_2. Nevertheless, when we combine E_1 and E_2, they together most-favor Dani's being agnostic.

How can this be? This is only possible, I believe, if Dani's reason for believing p (and her reason for believing $\neg p$) also favors, *to some degree*, being agnostic towards p. If Dani's reason for believing p somewhat favors being agnostic, and her reason for believing $\neg p$ somewhat favors being agnostic, then when these reasons are combined, they can overall most-favor her being agnostic.

To say that Dani's reason for believing p also favors, to some degree, being agnostic towards p, two things must be true:

1. Reasons favor sets of responses.[14]

[14]By "response", I mean the various kinds of things which reasons can be considerations in favor of. This includes, at least, actions, beliefs, and other attitudes. I borrow this terminology and much of the formalization

2. Favored responses are not necessarily favored to the same degree.

If (1) and (2) are true, we can quickly see how a reason for believing *p*, for example, can also (to some degree) favor being agnostic, and how when competing epistemic reasons are combined, they can end up most-favoring being agnostic towards *p*. Let's call the account of how reasons balance (which I am about to offer) which accepts (1) and (2) the *Structural Account*.

In defense of (1): though we generally speak of reasons favoring *an* action, or *an* attitude – such as when Scanlon (1998) says to be a reason is to "count in favor of" an action or attitude – it is reasonable to think that reasons instead favor sets of actions, or sets of attitudes. For example, the fact that I owe you $5 favors my paying you back. In virtue of this, I have a reason to pay you back. However, it seems that the thing being favored is not precisely an action. *Paying you back* is not a specific or individual action. It encompasses a whole group of actions, which count as actions of the sort paying you back. I could hand you a $5 bill. Or I could mail you $5. Or I could send you $5 digitally. And so on. These are all different actions, but they are all ways of paying you back. So, the fact that I owe you money favors all of the different actions which count as this type of action, or, as I have put it, it favors a *set of actions*.[15]

In defense of (2): often, all of the different ways I might do something (all of the actions in a favored set) are favored by the reason I have for doing them *to the same degree*. If, for example, I offer you money to raise a hand, then all of the different ways of raising a hand are equally rational with respect to the reason you have for acting (namely, the fact that you'll earn money). So, all of the actions in this set are favored to the same degree.[16] But not all favored sets are like this. Sometimes, members of a favored set can be favored to different degrees. Think again of the case where I owe you $5. There is a compelling sense in which giving you $4.99, rather than the full $5, is favored to *some* degree by the reason I have. It is certainly favored less than my giving you the full $5 – but it is clear that by giving you $4.99, I have at least partially responded to my reason. Furthermore, paying you back $4.99 is clearly better than only paying you $3, which is better than $1, and so on. It is hard to see how we can explain this without saying that there is *some degree* to which doing each of these things is favored by my reason for paying you back.

in this section from Skorupski (2010).

[15]Thank you to Jonathan Dancy for helping me see this point. See Dancy (2018, 31) for discussion of this example. Dancy calls this the "Prichard Point", in reference to Prichard (1932).

[16]This is separate from the claim that you have equal *overall* reason for doing any member of the favored set. In order to make a claim about overall reasons for engaging in some action, we must appeal to the total set of reasons that you have. The fact that I offer you money provides just one reason. But some additional considerations, like the fact that your left arm is hurting, can provide further reasons to weigh in your overall deliberation. When combined, these two reasons likely support raising your right hand, since this is the best way for you to earn some money without experiencing pain. But this is only in virtue of the combination of multiple reasons. My claim, then, is that the singular fact that I have offered you money to raise a hand, *in vacuo,* equally favors all the different ways you might do so.

Epistemic reasons also appear to often have this structure. Imagine that I have some evidence for *p,* which gives me epistemic reason for believing that *p*. It seems the most rational response to this reason (on its own) is to believe *p*. But we also want to say that being agnostic in light of such a reason is *more* rational than believing ¬*p*. This can only be explained if our reason for believing *p* also favors, to some degree, being agnostic whether *p*. So, we must describe being agnostic about *p* as being "in the set" favored by my epistemic reason – even if it is not favored to the same degree as believing that *p*.

If we accept (1), that reasons favor sets of responses, and (2), that the members of these sets are not necessarily favored to the same degree, then we can model a case of two balanced epistemic reasons in the following way. First, we'll say a *favoring function* describes how a fact favors a response for an agent. The function takes in an *f, a,* and *r* and outputs *d*, where *f* is a fact, *a* is an agent, *r* is a response, and *d* is the degree favored: *Favoring(f, a, r) = d.*

We can in turn define the *favored set*, denoted by *S*, as the set of responses which are favored to a degree greater than zero by *f* for *a*, as determined by the favoring function: *S* = {(r, d) | *Favoring(f, a, r) = d* and *d* > 0}.

Finally, we can define a *reason* as the relationship between a fact, an agent, and the favored set.[17] This can be written as: *Reason(f, a)* = S.

How does this way of representing reasons explain how competing reasons for *p* and for ¬*p* can most-favor being agnostic? Well, if being agnostic is a response in the favored set for each reason, then the two favored sets *intersect*. Being agnostic, as it falls within this intersection, is thus "double-favored". Take, for example, a reason for believing *p*, R_1, and a reason for believing ¬*p*, R_2. Say that the favored set for R_1, S_1, looks like this: S_1 = {(agnostic, 30), (believe, 50)} and S_2, the favored set for R_2, looks like this: S_2 = {(agnostic, 30), (disbelieve, 50)}. When R_1 and R_2 are combined to generate an overall verdict, we find that each response is favored to the following degree:

Disbelieve: 50
Agnostic: 60
Believe: 50

Since being agnostic is, overall, the uniquely-most-favored response to R_1 and R_2, being agnostic is the rational response for the agent. This explains our initial observation that when agents have balanced evidence for *p* and for ¬*p*, they ought to be agnostic whether *p* (and agnostic whether ¬*p*). Of course, what I have just provided is merely a toy model. But it demonstrates how competing reasons can combine to favor an intermediate response, given the assumptions that reasons favor sets of responses, and that members of these sets can be

[17] Note that there is disagreement about how many places the reason relation has. Skorupski (2010) takes it to have five places (fact, time, strength, person, action/belief), though it is generally agreed that at least three places are involved: facts, agents, and beliefs/actions. Nothing in my paper hangs on this.

favored to varying degrees.[18] I have argued these assumptions are independently plausible, and so I believe the Structural Account gives us the resources for resolving our puzzle.

The cases **Party or Library** and **Shy Friends** are now easy to explain. In **Party or Library**, Dani has some reason for believing that Charlie will show up to the party, and equal reason for thinking that Charlie will not show up to the party. We'd like to say that, in fact, Dani has *most* reason for being agnostic about whether Charlie will show up to the party. Now we can see why this is so: each of Dani's reasons, despite individually most-favoring something else, also favors to some degree being agnostic. When we combine all of her reasons, they overall most-favor her being agnostic. Enzo, in **Shy Friends,** is in a nearly-identical situation with regard to his reasons. He has a reason for sitting in chair #1, a reason for sitting in chair #2, and no other reasons. But, in combination, they most favor sitting in chair #3. So, neither Enzo nor Dani is experiencing a true case of balanced reasons.

Charlie, on the other hand, in **Party or Library**, does appear to be in a genuine balancing case. Each of his reasons individually most-favor going to the library or going to the party. When these reasons are combined, it turns out that no response is uniquely most-favored – going to the library and going to the party are truly balanced. This might be because there is no way of going to the library which also counts as a way of going to the party. More formally, the favored sets do not intersect. Since there is no uniquely most-favored response, Charlie can, with rational permissibility, simply pick between going to the library and going to the party.

In sum, the answer to our puzzle is this: a reason for some first thing, and a reason for some second thing, can, *in combination*, most-favor some third thing. This is underwritten by two plausible features of reasons which I have called the Structural Account. In the next section, I address a potential objection to this way of resolving the puzzle.

5 An Objection

I have been arguing – against the Balancing Assumption – that our best explanation for observations about the behavior of reasons includes the denial of prohibitively-balancing reasons. I have tried to do so in a way that introduces no new essential differences between practical and epistemic reasons. However, there is a natural objection to be made. All that I have said so far seems to commit us to something like the following conditional:

[18] Though this treatment has been in terms of full belief/agnosticism/disbelief, I am optimistic that such an account could also be put in terms of degrees of belief. The Brier Score, for example, is one way we could specify the values of epistemic favoring functions. For example, If Dani's evidence equally favors p and $\neg p$, adopting a credence of 1 (or 0) in p would maximize inaccuracy with an expected square error of .5. So, assuming Dani ought to maximize accuracy, she instead ought to adopt a credence of .5 in p, so that her expected square error is only .25. This would explain why a credence of .5 is most-favored by Dani's combined reasons. I think such a way of bridging the gap between reasons-talk and credal epistemology deserves more attention.

> **Epistemic Permission**: If epistemic reasons ever genuinely balance, then they balance permissively.

Nothing about my rejection of the Balancing Assumption or use of the Structural Account precludes the possibility of genuinely balanced epistemic reasons. Therefore, this account seems to predict a form of *permissivism*, the negation of the thesis that given one's total evidence about p, there is a unique doxastic attitude that one rationally ought to have towards p.[19] Put bluntly, permissivism is a controversial thesis.

Here is the objection, sharply: the foregoing account seems to predict that there are cases of genuinely permissively-balanced epistemic reasons. But we observe no such cases. So, the foregoing account is false.

There are two ways of responding to this objection, corresponding to each element of the **Epistemic Permission** conditional. The first way of responding is to deny the antecedent: claim that epistemic reasons are *never* in a state of balance. This would make **Epistemic Permission** vacuously true. The second way of responding is to accept the consequent: to argue that we *do* observe cases of permissively-balanced epistemic reasons. As I alluded to in the introduction, the former response will likely be convincing to evidentialists but not pragmatists, and the latter response will likely be convincing to pragmatists but not evidentialists. Thankfully, the success of what I have been arguing for in this essay does not depend on getting evidentialists and pragmatists to agree on anything more than the rejection of the Balancing Assumption. Each party should draw their own lessons – lessons I will now attempt to spell out.

I'll start by reviewing potential cases of genuinely balanced epistemic reasons. One example is due to Roeber (2016), who suggests that there will be cases in which it is permissible to believe p and at the same time permissible to be agnostic towards p:

> **Eye Exam** Imagine looking at an eye chart. At the top of the chart, there is a very large, perfectly clear 'E.' At the bottom, there's a bunch of letters, but they look like dots. The transition from the perfectly clear 'E' at the top of the chart to the blurry dots at the bottom is very gradual. What happens as you move from the top of the chart to the bottom?

It seems, at the top of the chart, that you ought to believe the letter is 'E'. Moving down a step, the subsequent letter is less-so but still quite clear, so you ought to believe what you see. And so on, until it's too blurry to be sure. Conversely, starting from the bottom, you ought to be agnostic about what letter you are seeing, since you can merely see blurry dots. As

[19] See White (2005), Schoenfield (2012), Kelly (2013) for more on permissivism. The permissivism which appears to be entailed by my account is an intra-personal version of permissivism. By this I mean that a particular epistemic agent, in the right situation, could have rational permission to choose between different doxastic attitudes. The debate about permissivism often instead focuses on inter-personal permissivism, which is the question of whether two different agents with the same body of evidence could each rationally have different doxastic attitudes. It seems then that the permissivism which may be entailed by my account is an even stronger version of an already controversial thesis.

you move upward, you become more confident in which letter you are seeing. Eventually, Roeber suggests, you reach a point in the middle of the chart where your reasons are in perfect balance, and its thus permissible to believe or to be agnostic about whether you are seeing the letter 'E'. If Roeber is right, then this is a genuine case of epistemic reasons being balanced – and they balance permissively. This would satisfy **Epistemic Permission.**

A second case is inspired by Moss (2011):

> **Evil Scientist** An evil scientist tells you he is going to perform an operation to change your attitude towards a certain proposition p, about which you are currently committedly agnostic. The scientist gives you a choice: after the operation, you may either believe p or believe $\neg p$.

Here we can assume that you have preexisting epistemic reasons, which, when combined, uniquely recommend being agnostic about p. But what if you *cannot* be agnostic, like in **Evil Scientist**? Then it seems your epistemic reasons no longer uniquely most-favor any particular attitude. In the language we have developed in previous sections, we can say that two responses (believe p and believe $\neg p$) are equally most-favored by your combined reasons. Since there is equal epistemic reason for you to believe p as there is for you to believe $\neg p$; you should just pick one. In other words, your epistemic reasons are balanced permissively. This accords with the intuition that there's nothing epistemically irrational about believing p or believing $\neg p$ in such a case.[20]

A third suggestion. In the same way we have suggested that there are many ways of doing certain actions, each of which is equally-favored by certain individual facts, perhaps we can also say that there are many ways of *believing* which are equally-favored by our reasons. This point is made by Dancy (2018): "believing is a response, and as such it will, just like an action, have various properties about which the considerations that favour it are silent. For instance, it can be done enthusiastically or reluctantly, it can be done sooner or later—and so on." Let's say Dancy is right about this. Then, if our epistemic reasons happen to uniquely favor believing p (rather than being agnostic or disbelieving p), then these reasons may still be silent on whether to do so enthusiastically or reluctantly. So, we might in fact think that our reasons *are* permissively balanced, between believing p in one way and believing p in another. This would be a case where we can permissively choose between our equally-most-favored options.

I have just suggested three potentially genuine cases of permissively balanced epistemic reasons. I suspect, however, that none of these considerations will be convincing to committed evidentialists. Such philosophers, for example, may deny that Roeber's case is truly permissive, deny that there are epistemic reasons for action, and argue that Dancy's way of individuating reasons is misleading.

[20] In this example, we might think that the response is an action rather than an attitude. However, it seems the reasons are still epistemic ones. This suggests that there could be epistemic reasons for action, and that such epistemic reasons can end up in a genuine state of balance. See Booth (2006) for a defense of epistemic reasons for action.

What if we instead go evidentialists' likely-preferred route of denying that epistemic reasons are ever genuinely balanced? The problem with going in this direction is that it appears to once again introduce essential differences between practical and epistemic reasons. One might think as follows: "I accept that epistemic reasons don't balance prohibitively, so in this sense there is no difference between practical and epistemic balancing. But, also, epistemic reasons never balance, and practical reasons at least sometimes do, so nevertheless we still find a major difference". The worry, then, is that we have merely redescribed our problem.

I think this worry is misplaced. The Balancing Assumption takes the balancing behavior of reasons to be a property of reasons themselves. By shifting away from this sort of explanation, we move towards explaining balancing phenomena in terms of the situations in which reasons are salient. Evidentialists may be right, in the end, that epistemic situations are entirely divorced from the practical ones – but this is no longer a claim about epistemic and practical reasons *in-themselves*.

Additionally, (as we have seen) the Balancing Assumption is often used as a premise in arguments against unification and commensurability. So, what happens if we remove the Balancing Assumption from such arguments, and substitute in the fact that epistemic reasons never genuinely balance? If we have merely redescribed our problem, then this premise-substitution should not affect such arguments. In §6, I'll argue that engaging in such premise-substitution *does* impact their conclusions. If I am right, then this is strong evidence that we have not merely given our problem a fresh coat of descriptive paint.

6 Removing the Balancing Assumption

Let's first return to the challenge to Schroeder's (2015) unifying account of what makes reasons sufficient. Schroeder calls his unified thesis about reason-sufficiency *Sufficiency as Balance*:

> **Sufficiency as Balance**: A set S of reasons for X in favor of A is sufficient just in case for each alternative B to A, S is at least as weighty as the set of all of the reasons for X in favor of B.

If the Balancing Assumption were true, then it is obvious why Sufficiency as Balance would not explain the balancing behavior of epistemic reasons. Take our case **Will He Attend**. Dani has exactly equal epistemic reason for believing that Charlie will show up to the party as she does for believing that Charlie will not show up to the party. According to Sufficiency as Balance, a set of reasons is sufficient just in case it is at least as weighty as any other competing set of reasons. So, according to Sufficiency as Balance, Dani has sufficient reason for believing Charlie will show up, and for believing he will not show up. But if the Balancing Assumption is true, then Dani's balanced epistemic reasons prohibit her from believing Charlie will show up, and prohibit her from believing Charlie will not show up. So, the Balancing Assumption and Sufficiency as Balance are incompatible.

I have argued, however, that the Balancing Assumption is false. In most cases where it looks like one's epistemic reasons are balanced, like in **Will He Attend,** there is really only one unique attitude which is uniquely favored by their reasons. Therefore, in Dani's case, her set of reasons for being agnostic are (to put it in Schroeder's terms) at least as weighty as any other competing set of reasons, and thus she has sufficient reason for being agnostic. Furthermore, if it is true that epistemic reasons are never genuinely balanced, then Sufficiency as Balance always gets the right result. So, Schroeder's unified theory of sufficiency is no longer undermined by these kinds of considerations.

What of commensurability? Berker (2018), an evidentialist, uses the Balancing Assumption to argue that it's impossible to rationally weigh practical reasons against epistemic reasons. Berker argues that if the Balancing Assumption is true, and practical and epistemic reasons balance in essentially different ways, then this leads to highly counter-intuitive results when attempting to weigh one kind against the other. Of course, we suspect that the Balancing Assumption is false. If we have merely redescribed the problem, then Berker's argument ought to still work. Let's check.

First, assume that epistemic reasons *do* balance prohibitively. Suppose that an agent, Fiona, has extremely strong, and exactly equal, practical reasons for being agnostic whether p, and for disbelieving p. If these practical reasons exhaust all of Fiona's reasons (she has no epistemic reasons), then she has all-things-considered permission to either be agnostic whether p or to disbelieve p. This fits our usual intuitions – if Fiona has been offered \$10,000 to disbelieve p, and \$10,000 to be agnostic whether p, and has no other practical reasons, and no epistemic reasons for or against p, then she's permitted to disbelieve p or be agnostic whether p.

If Fiona tries to weigh these existing practical reasons against incoming epistemic reasons for or against p, Berker continues, things will not end well. Say that at t_1 Fiona obtains exactly equal epistemic reasons for believing p as for believing $\neg p$. Her epistemic reasons, considered on their own, appear to prohibit her from believing either p or $\neg p$, so her purely epistemic verdict is to be agnostic on whether p. Being agnostic is also permitted by her practical reasons, so it seems that Fiona all-things-considered should be agnostic whether p.

Verdict	Practical Reasons	Epistemic Reasons	All-things-considered
Disbelieve	Permitted	Prohibited	Prohibited
Agnostic	Permitted	Permitted	**Permitted**
Believe	Prohibited	Prohibited	Prohibited

Imagine at a later time, t_2, Fiona gains new strong epistemic reasons in favor of believing that p. Based on her epistemic reasons alone, it seems she ought to believe p. But assume that her practical reasons (which pay out \$10,000 for disbelief or being agnostic) are still

much weightier than her epistemic reasons.[21] It seems rational, in this case, for Fiona to respect her practical reasons *over* her epistemic reasons, since they outweigh her epistemic ones. When she therefore weighs all of her reasons, practical and epistemic, Fiona gets the all-things-considered verdict that she is permitted to disbelieve p, and permitted to be agnostic whether p, since these are the two responses which are equally most-favored.

Verdict	Practical Reasons	Epistemic Reasons	All-things-considered
Disbelieve	Permitted	Prohibited	**Permitted**
Agnostic	Permitted	Prohibited	**Permitted**
Believe	Prohibited	Permitted	Prohibited

However, when we compare this verdict to the verdict at t_1, the change in Fiona's permissive doxastic states seems odd, or inconsistent. At t_1, Fiona all-things-considered should have suspended judgement. Then, at t_2, she gained conclusive epistemic reasons for p. However, by obtaining conclusive epistemic reasons for p, Fiona gained permission to *disbelieve* p – something we thought she was *prohibited* (to use the language of the Balancing Assumption) from doing when she had less evidence for p. In other words, it appears that Fiona was able to "unlock" permission to disbelieve p by gaining more evidence for p. This looks strange. Berker suggests that the source of this problem is that we're trying to weigh practical and epistemic reasons against each other, which cannot be done. He argues that any procedure by which we might try to weigh practical and epistemic reasons against each other will result in highly counter-intuitive verdicts like this, because of the different ways practical and epistemic reasons balance.

But Berker's argument relies on the Balancing Assumption. Let's consider what happens in this case if we replace the Balancing Assumption with our Structural Account of how reasons balance (and the thesis that epistemic reasons are never balanced).

Suppose, as we did before, that Fiona has extremely strong practical reason to either be agnostic whether p, or to disbelieve p. At t_1, it is also the case that her purely epistemic reasons combine to most-favor being agnostic whether p. Since this purely epistemic verdict does not conflict with her very strong practical reasons, it seems she should all-things-considered be agnostic whether p. So far, so good.

Verdict	Practical Reasons	Epistemic Reasons	All-things-considered
Disbelieve	Equally Most-Favored	Least Favored	Less Favored
Agnostic	Equally Most-Favored	Most Favored	**Most Favored**
Believe	Not Favored	Least Favored	Less Favored

[21] To make this intuitive, readers should feel free to adjust the payout number upwards (and the innate importance of p downwards) until it's clear to them that the practical reasons outweigh the epistemic ones.

As before, at t_2, Fiona gains new epistemic reasons to believe that p. She still has overwhelming practical reason to either disbelieve that p or be agnostic whether p. Her new epistemic reasons do not give Fiona enough overall reason to believe that p, since it is outweighed by her practical reasons (by our stipulation). However, Fiona can still recognize that her epistemic reasons favor being agnostic *more* than they favor disbelief. This means that, of the attitudes left open by her practical reasons, being agnostic is the one most favored.

Verdict	Practical Reasons	Epistemic Reasons	All-things-considered
Disbelieve	Equally Most-Favored	Least Favored	Second-Most Favored
Agnostic	Equally Most-Favored	Second-Most Favored	**Most Favored**
Believe	Not Favored	Most Favored	Least Favored

Therefore, when Fiona gains epistemic reasons for believing p at t_2, the most rational attitude for her is still to be agnostic. Unlike when we attempted to balance prohibitive and permissive reasons, Fiona does not "unlock" rational permission to disbelieve p by gaining epistemic reasons for believing p. There is nothing counter-intuitive about this - so it seems that without the Balancing Assumption, Berker's argument no longer goes through.

7 Conclusion

Even if we believe epistemic reasons never genuinely balance, I think the considerations of the previous section show that we can safely dispel with worries of "mere redescription". As for the argument overall, the lesson to be drawn is that the Balancing Assumption does not adequately explain our observations of practical and epistemic balancing. I have argued that the Structural Account does a better job. It does so by showing how it is possible for competing reasons to *combine* to most-favor actions/attitudes that neither reason most-favors on its own. This, in turn, explains how apparent cases of prohibitive balancing are not really cases of balancing at all – but merely misdescribed cases where a unique response is most-favored.

Where does this leave us? Pragmatists, I suspect, will feel that this way of understanding how reasons balance neatly resolves our puzzles. Evidentialists, on the other hand, who are likely not convinced by the kinds of permissive-balancing cases in §5, are left with a new question: what is the deeper explanation for the fact that epistemic reasons always favor one unique response?

References

10 Berker, Selim. 2018. "A Combinatorial Argument Against Practical Reasons for Belief." *Analytic Philosophy* 59 (4): 427–70. https://doi.org/10.1111/phib.12140.

Booth, Anthony Robert. 2006. "Can There Be Epistemic Reasons for Action?" *Grazer Philosophische Studien* 73 (1): 133–44. https://doi.org/10.1163/18756735-073001008.

Brunero, John. 2022. "Practical Reasons, Theoretical Reasons, and Permissive and Prohibitive Balancing." *Synthese* 200 (2): 1–23. https://doi.org/10.1007/s11229-022-03628-8.

Cohen, Stewart. 2016. "Reasons to Believe and Reasons to Act." *Episteme* 13 (4): 427–38. https://doi.org/10.1017/epi.2016.22.

Dancy, Jonathan. 2004. *Ethics Without Principles*. New York: Oxford University Press.

———. 2018. *Practical Shape: A Theory of Practical Reasoning*. Oxford: Oxford University Press.

Feldman, Richard. 2006. "Epistemological Puzzles About Disagreement." In *Epistemology Futures*, edited by Stephen Cade Hetherington, 216–36. Oxford University Press.

Friedman, Jane. 2013. "Suspended Judgment." *Philosophical Studies* 162 (2): 165–81. https://doi.org/10.1007/s11098-011-9753-y.

Harman, Gilbert. 2004. "Practical Aspects of Theoretical Reasoning." In *The Oxford Handbook of Rationality*, edited by Alfred R. Mele and Piers Rawling, 45–56. Oxford University Press.

Hieronymi, Pamela. 2005. "The Wrong Kind of Reason." *Journal of Philosophy* 102 (9): 437–57. https://doi.org/10.5840/jphil2005102933.

Horty, John F. 2012. *Reasons as Defaults*. Oxford, England: Oxford University Press USA.

Kelly, Thomas. 2002. "The Rationality of Belief and Other Propositional Attitudes." *Philosophical Studies* 110 (2): 163–96. https://doi.org/10.1023/a:1020212716425.

———. 2013. "Evidence Can Be Permissive." In *Contemporary Debates in Epistemology*, edited by Matthias Steup and John Turri, 298. Blackwell.

McGrath, Matthew. 2021. "Being Neutral: Agnosticism, Inquiry and the Suspension of Judgment." *Noûs* 55 (2): 463–84. https://doi.org/10.1111/nous.12323.

Moss, Sarah. 2011. "Scoring Rules and Epistemic Compromise." *Mind* 120 (480): 1053–69. https://doi.org/10.1093/mind/fzs007.

Parfit, Derek. 2011. *On What Matters: Volume One*. OUP Oxford.

Raz, Joseph. 1999. *Engaging Reason: On the Theory of Value and Action*. Oxford, GB: Oxford University Press UK.

Reisner, Andrew. 2008. "Weighing Pragmatic and Evidential Reasons for Belief." *Philosophical Studies* 138 (1): 17–27. https://doi.org/10.1007/s11098-006-0007-3.

———. 2009. "The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem." *Philosophical Studies* 145 (2): 257–72. https://doi.org/10.1007/s11098-008-9222-4.

Rinard, Susanna. 2017. "No Exception for Belief." *Philosophy and Phenomenological Research* 94 (1): 121–43. https://doi.org/10.1111/phpr.12229.

———. 2019. "Equal Treatment for Belief." *Philosophical Studies* 176 (7): 1923–50. https://doi.org/10.1007/s11098-018-1104-9.

Roeber, Blake. 2016. "Reasons to Not Believe (and Reasons to Act)." *Episteme* 13 (4):

439–48. https://doi.org/10.1017/epi.2016.23.

Scanlon, Thomas, ed. 1998. *What We Owe to Each Other*. Cambridge: Harvard University Press.

———. 2014. *Being Realistic About Reasons*. Oxford: Oxford University Press.

Schoenfield, Miriam. 2012. "Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief." *Noûs* 48 (2): 193–218. https://doi.org/10.1111/nous.12006.

Schroeder, Mark. 2015. "What Makes Reasons Sufficient?" *American Philosophical Quarterly* 52 (2): 159–70.

———. 2021. *Reasons First*. Oxford University Press.

Skorupski, John. 2010. *The Domain of Reasons*. Oxford: Oxford University Press.

White, Roger. 2005. "Epistemic Permissiveness." *Philosophical Perspectives* 19 (1): 445–59. https://doi.org/10.1111/j.1520-8583.2005.00069.x.