

Dynamic Permissivism

Abelard Podgorski

Introduction

There has been considerable recent philosophical debate over a thesis called *epistemic permissivism*. According to the permissivist, it is possible for two agents to have the exact same total body of evidence and yet differ in their doxastic attitudes towards some proposition, without either being irrational. Although the camp of permissive views is large and diverse, it has found its fair share of criticism, most notably in a seminal paper by Roger White (2005) and the vigorous debate it spawned.

However, I will argue, not enough attention has been paid to the distinction between different ways in which permissivism might be true. In particular, there is one type of permissive view which the contemporary debate has completely ignored. This view permits differences in belief between evidential equals by rejecting *state-oriented* norms connecting an agent's rational behavior to their evidence in favor of *dynamic* norms governing their processes of consideration and belief-update.

In this paper, I will present a taxonomy of forms of epistemic permissivism framed as the upshot of different ways one might respond to a basic argument against the view culled from White. I will discuss briefly some of the commitments these views incur and the resources they have to answer challenges from their critics. Then I'll show how the dynamic strategy of rejecting state-oriented norms on belief opens the door to a new and particularly attractive kind of permissivism. In doing so, I will give an independently plausible sketch of the kinds of norms such a picture would include, discussing their implications for rational behavior in general and the permissivist dispute in particular. This view, I will argue, is especially well-placed to answer worries that have been raised against traditional permissivist views.

1.1 – Permissivism and Uniqueness

Epistemic permissivism, roughly speaking, is the claim that agents with the same evidence can reasonably disagree (the denial of permissivism has its own name, *uniqueness*). More precisely:

Permissivism: It is possible for two agents (or the same agent in different possible worlds), to a) have the same body of total evidence E, b) hold different doxastic attitudes towards proposition P, and c) exhibit no irrationality.

I will not take a stand here on exactly what constitutes our total evidence, thinking of it as, minimally, some set of mental states like belief or knowledge (or the content of such states) which are in some broad sense internally accessible to us. To be fair to uniqueness, I will assume that our evidence is relatively extensive, including things that are only a passive part of our psychology at the time. Taking an overly narrow view of evidence, such as limiting it to those things of which we are at the time immediately consciously aware, would render Uniqueness implausible on its face, since our evidence at any given time would be far too sparse to settle anything like the multitude of beliefs ordinary agents carry with them.

Although views with implications concerning permissivism have been around for a long time, the contemporary debate on the topic has as its focal point a paper by Roger White (2005) in which he presents a series of arguments against the view. Since then, both sides have found their fair share of advocates – besides White himself, Feldman (2007), Christensen (2007), and Hedden (2015) come out in favor of uniqueness, while Meacham (2013), Schoenfield (2014), and Kelly (2013) defend versions of permissivism.

I will, for the purposes of this paper, be primarily considering permissivism as a view concerning *full belief*, as opposed to other doxastic attitudes like *credences*, which are foremost in the mind of those

disputants in the literature of a bayesian persuasion.¹ This is primarily for the purposes of unburdened presentation, and the dialectic over permissivism in these other attitudes is similar, though the reader is advised some kinds of permissivism may be more or less plausible when applied to them.

Secondly, although permissivism can take a *moderate* form, according to which some given evidence may allow only slack between believing a proposition and withholding judgment on it, I will be concerned with a more *extreme* kind of permissivism, which holds that slack exists between believing a proposition and believing the negation of that proposition. I have a particular interest in the prospects of this more radical view since it is a consequence of the dynamic approach I favor, and many proponents of permissivism in the literature do endorse it in its stronger form. Still, although my discussion will be framed around the debate over the extreme view, the taxonomy I will develop equally well describes options for a moderate approach.²

1.2 – Why Permissivism?

Permissivism as such is rarely argued for in depth. It is sometimes claimed to have the force of intuition behind it. Gideon Rosen (2001) suggests that “It should be obvious that reasonable people can disagree, even when confronted with the same body of evidence. When a jury or a court is divided in a difficult case, the mere fact of disagreement does not mean that someone is being unreasonable.” (p. 71) In light of subsequent controversy, the claim of obviousness seems premature, though the general sentiment, echoed in Schoenfield (2014), retains some force; given the ubiquity of disagreement between intelligent and similarly informed people, rejecting permissivism implies that irrationality is more widespread than we generally think.

Much of the rest of the support for permissivism comes from the fact that it follows from a number of epistemic views philosophers have found independently attractive. White (2005), for example,

1 Meacham (2013) defends the permissive form of Bayesianism, while Christensen (2007) argues for impermissivism.

2 For a defense centered on the more moderate versions of permissivism, see Kelly (2014). Also, Mark Nelson (2010) argues for a moderate view on which withholding judgment is *always* permissible.

lists conservatism, subjective bayesianism, and reflective equilibrium as established classes of views committed to permissivism, and Ballantyne/Coffman (2011) argue that uniqueness rules out several plausible pairings of views about evidence and rationality internalism – possible stances on whether internal duplicates can differ in their evidence and their rationality respectively. If attractive models of rational belief are committed to permissivism, then that is a presumption favoring the view.³

I will not be breaking this dialectical pattern here. I take it for granted that there is considerable variety in the doxastic attitudes of ordinary people with roughly similar evidence, that we do not naturally attribute to irrationality on the part of one or more of the parties, and that it is a virtue of permissivism that it provides a simple face-value explanation of this fact. And though it is beyond the scope of this paper to argue, I believe the thesis of epistemic conservatism, that the fact an agent already believes something can make a difference to whether it is rational for them to believe it going forward, is independently plausible for a number of reasons.⁴ Such a view entails that there are at least two different doxastic attitudes that might be available to an agent with certain given evidence – the attitude that it is rational for them to have going forward if they do not already believe P, and the attitude that is rational if they do. Any attractions of conservatism, then, are ipso facto attractions of permissivism.

I don't take these meager considerations to be sufficient to motivate those committed to uniqueness to abandon their views. My task in this paper is to explore ways in which permissivism *could* be true and assess their relative advantages, as the thesis can be achieved in a number of ways, the distinctions between which have not been fully appreciated, and to welcome a new view into the family that I believe stands on its own and may ameliorate skepticism of the permissivist outlook. Before I present the dynamic form of permissivism I prefer, I will canvass the more common varieties and the difficulties they face. The differences between the views will stand out more clearly, however, if we introduce them in light of a basic

3 In particular, they suggest that uniqueness precludes both being an internalist about rationality and an externalist about evidence, and being an externalist about rationality and an internalist about evidence.

4 Some motivations for conservatism can be found in Harman (1986), McGrath (2007).

challenge to the general permissivist project.

2.1 – The Evidence-Pointing Argument

The canonical attack on permissivism comes from Roger White (2005). He presents and elaborates his objections in several ways, but it will be instructive to first consider a version of what he calls his “quick and dirty” argument against the view (p. 447). Following Schoenfield (2014), we may call it the “Evidence-Pointing Argument”, and it runs as follows:

- (1) ***State-Oriented Evidential Support***: It is rational for an agent to believe that P only if their total evidence supports P.
- (2) ***The Evidence Points One Way***: Any given set of evidence may support P or it may support not-P, but it cannot support both.
- (3) Therefore, there is no body of evidence such that sometimes an agent with that total evidence may rationally believe P and sometimes an agent with that evidence may rationally believe not-P.

Though White takes the evidence-pointing argument to be sound, he recognizes that it may not be fully dialectically adequate. Nevertheless, it is a helpful tool because we will find, in different responses to this argument, the seeds for distinctive species of permissivism.

2.2 – Options Permissivism

One way to respond to the evidence-pointing argument is to reject premise (2) and insist that a single set of evidence can, at the same time, in the same sense, and for the same person support both P and not-P. When one considers one's evidence, on this view, multiple options present themselves as equally supported. I call this approach *options-permissivism* because it allows that a single agent with given evidence

has, in some cases, something like an epistemic *choice* between believing P and believing not-P. It allows *intrapersonal* slack between evidence and belief. Though this extreme, full-belief version of options-permissivism is unpopular for reasons that will become clear, Kelly (2010, 2013, p. 300) expresses sympathy with a version of this kind of view applied to credences.

2.3 – Background-Relative Permissivism

There are two sibling permissive views which take a very different approach. In response to the evidence-pointing argument, these views do not reject 2 outright, but insist that it has a missing parameter. The support relation, it is claimed, is not a mere two-place relation between a set of evidence and a belief or proposition. It is a three-place relation between a set of evidence, a belief or proposition, and some *background* condition of the agent.

On one spelling-out of this background condition, evidence supports a proposition only relative to a certain epistemic *starting point*. The cleanest example of such a view is the sort of subjective bayesian picture defended, in the context of the permissivism debate, by Meacham (2013). According to the standard bayesian picture, rational epistemic behavior is determined by two things – one’s evidence, and one’s prior probability function. The prior operates as a kind of essential background against which evidence does its work. Given a particular prior, there is only one way an agent can respond to her evidence. So this view does not tolerate intrapersonal slack in the way the options-permissivist does. The view secures permissivism by claiming that there are many equally rational priors that can be had by different agents – agents with the same evidence but with different priors will rationally come to different conclusions.

A second strategy of this type is defended by Schoenfield (2014) and takes evidential support to be relative some set of *epistemic standards*. Much like priors, the epistemic standards determine how the agent is to respond to evidence, and much like priors, many different epistemic standards are equally rational. Two

agents may, with full rationality, have the same evidence and different beliefs due to a difference in their epistemic standards, while each agent's standards only allow one response *for that person*.

These views both claim, then, that to be true, (1) must be read as claiming “it is rational for an agent to believe P only if their total evidence supports P *relative to their own background*. And to be true, (2) must be read as something like “Any given set of evidence may support P *relative to background Q* or it may support not-P *relative to the same background*, but not both”. But on these readings, (3) no longer follows from (1) and (2). So, the evidence-pointing argument is blocked.

2.4 – Non-evidential Permissivism

A third way of responding to the evidence-pointing argument involves rejecting premise 1. This premise is a norm that is both *state-oriented* and *evidential*. It is state-oriented because it judges an agent’s rationality in virtue of features of their states of mind at individual times – in this case their beliefs. It is evidential because the constraint it places on the agent is determined entirely by the agent’s evidence. Because there surely is *some* norm that connects an agent’s beliefs to the considerations that support those beliefs, and because evidence is paradigmatically the sort of thing that bears on belief, state-oriented evidential support has a great deal of initial plausibility. But some deny that evidential considerations exhaust the considerations that bear on belief. The strategy of what I’ll call *non-evidential permissivism* is to replace the evidential state-oriented support principle in (1) with a *non-evidential* state-oriented support principle, that is, to claim instead that the right requirement should be something like:

State-oriented Non-evidential Support: It is rational to believe P only if the evidence, *together with all the relevant non-evidential factors* supports P.

Replacing (1) with such a principle blocks the evidence-pointing argument. The evidence may point

only one way, the non-evidential permissivist will say, but the evidence *plus* other non-evidential factors will, in different cases, point in different ways.

Different versions of this view will take different stances on what the relevant non-evidential factors are. One possibility is that these are *pragmatic* reasons for belief. While this sort of view is unpopular, a moderate picture of this sort is defended by Schroeder (2012), who argues that there are pragmatic reasons *against* belief, which can lead agents with the same evidence to differ in whether they rationally believe or withhold on a proposition. An extreme permissivism in this vein would have to identify pragmatic factors that can make the difference between believing P and believing its negation.

A more common approach of this type we have mentioned already – standard *epistemic conservative* views are a kind of non-evidential permissivism that take the mere fact that one already believes in P to be a non-evidential epistemic factor in favor of P. Proponents of this view argue that the conservative view has several important virtues, from allowing us to manage the cognitive costs of changing our minds (Lycan 1988) to providing a solution to skepticism (Sklar 1975). Permissivism falls out of this view as a corollary, since it allows agents who differ in whether they already believed a proposition to differ in whether they rationally believe it going forward, even when their evidence is the same.

3.1 – An Evaluation

It is not my aim in this paper to refute these forms of permissivism. But to whet the reader's interest in my own special brand of the view, it will help to see why one might worry about traditional elaborations of the permissivist thesis and the commitments they incur.

Options-permissivism seems to be the sort of permissive view critics (White 2005, Christensen 2007, Feldman 2007) have in mind by default. But few of those engaged in active defense of permissivism commit to this view when it comes to a belief and its negation, and it is not difficult to see why. First of all, *The Evidence Points One Way* is extremely plausible on its face, and the options-permissivist doesn't seem

to have in easy reach any way to cushion the blow of rejecting it by, for instance, proposing a replacement, as the other forms of permissivism do. It's difficult to see what it could be for evidence to support some proposition except by *contrast* with the negation – to support that proposition *over its falsehood*. Each part of our evidence counts in favor of P, it is natural to think, to exactly the same extent it counts against \sim P. Evidential support, on this intuitive picture, works like a scale with the evidence for P piled on one side and the evidence for \sim P piled on the other. To say that the evidence could point both ways is to suggest that there are ways of placing weights on the scale that lift both ends at once. On this way of thinking, to claim that the evidence can on balance support both P and \sim P sounds incoherent.

Furthermore, the idea that we may be presented with an epistemic choice between believing P and believing \sim P is false to our own experience. When we form a belief that P, we do not do so with the background thought that the evidence for \sim P is just as compelling. Indeed, to form a belief like this is a borderline psychological impossibility. In cases where the evidence for P and the evidence for \sim P are equally compelling, the only rational response, and the one human agents actually manifest seems to be *withholding judgment*. Options permissivism's relatively radical commitment to intrapersonal slack does not, moreover, do much work in explaining the kinds of intuitions that motivate permissivism – which are typically thoughts about rational *disagreement*. So the considerable costs that attach to this feature of the view come with very little upside.

The idea that there might be epistemic options available to us is most plausible in borderline cases – where the evidence is right around the boundary of being good enough for a belief that P. And Kelly (2014) suggests that some given evidence may leave rationally open whether to have, say, 39% or 40% credence in some proposition. But these boundaries exist between believing something and withholding judgment, or between one credence and another, very nearby credence, not between believing something and believing its negation. So options-permissivism looks like it's poised, at best, to account for a *moderate* permissivism.

Background-relative views are in a somewhat better position. By recovering a plausible reconstrual of the highly intuitive premise (2), and because they are not committed to problematic intrapersonal slack, these views have a better response to the evidence-pointing argument than options-permissivism.

But one may worry that these views have something of a needle to thread. The strategy owes us the promise of some kind of principled explanation of the *extent* of slack in the background conditions. There are dangers in two directions. Consider the epistemic-standard permissivist view. If *any* epistemic standards are fair game, then no matter how bizarrely someone reacts to evidence (by, say, forming random or even contradictory beliefs), they will not be rationally criticizable provided they are following their own rules. Similarly, if someone's priors are allowed to be sufficiently extreme, their Bayesian responses to evidence will be unrecognizable as what we would ordinarily consider rational behavior. If, on the other hand, the range of acceptable standards or priors is very narrow, then the view may end up too moderate to explain the range of disagreement we find intuitively rational, and lose out on one of the motivations for a permissivist outlook. Whether any sensible, non-ad-hoc restriction of this sort is available is far from clear.

The non-evidential permissivist has their own set of concerns. It is generally thought, from reflection upon Pascal's Wager and similar cases, that pragmatic factors are reasons of the *wrong kind* to bear on a belief's epistemic rationality. Even some who think that there are such pragmatic reasons claim that they specifically concern *withholding* (See Schroeder 2012). Because on such views non-evidential factors can only make the rational difference between believing P and withholding on whether P, and never between believing P and believing not-P, they only justify permissivism in a moderate form. The other main representative from this category, the standard conservative view, on the other hand, faces several powerful challenges, threatening to license agents in bootstrapping themselves into rational beliefs merely by forming them, and objectionably treating an agent's own beliefs asymmetrically from those of others (Christensen 1994). In addition, any non-evidentialist view rejects a plausible claim we may call the *evidence*

restriction, the thesis that the only things bearing on beliefworthiness in a distinctly epistemic way are *evidential*, relevant in some way on the belief's likelihood of being true. The evidence restriction is both intuitive and theoretically attractive for the way it draws the boundary between reasons for belief of the right, genuinely epistemic, kind and the wrong kind, while reflecting what many take to be belief's fundamental aim of truth.

3.2 – The Arbitrariness Argument

There is also a general worry about permissivism arising from the thought that these views make our beliefs in some sense objectionably *arbitrary*. Although he does not give the argument in this form, the following is, I think, a distillation of a particularly compelling set of concerns White raises in developing his objection:

- (A) Any plausible form of (extreme) permissivism will allow that we can sometimes know we are in a permissive case and retain our belief.
- (B) If I believe I am in a permissive case, I must view any belief on the topic as arbitrary from the standpoint of truth.
- (C) Viewing my belief as arbitrary from the standpoint of truth makes it irrational for me to hold the belief.
- (D) So, it is impossible for me to rationally hold a belief when I know I am in a permissive case.
- (E) So, (extreme) permissivism is false.

One might reject (A) if one thinks that permissive cases only exist for those who are blind to them – that rational disagreement may exist only when the parties involved do not realize that it does. But this self-effacing permissivist response is quickly discarded in White's own discussion (pp. 450-451), and to my

knowledge no permissivist is happy to endorse it. This view has the odd implication that meta-information vindicating your present belief as rational can defeat that belief. Moreover, permissivists typically think the phenomenon of rational disagreement is fairly pervasive – indeed, as we have noted, the appeal to intuitions to that effect are part of the motivation for the view. There is a tension between the visibility this implies for permissive cases and the blindness rejecting (A) suggests we need for them to exist. (C) is also difficult to deny. It is arguably constitutive of belief as a mental type that it is fundamentally concerned with truth. (B), it seems, is the crucial premise. We may ask, then, how plausible it is given each of the forms of permissivism we have so far discussed.

Options Permissivism seems especially poorly positioned to deny B). After all, on that view, believing P and believing not P are both direct options for me, once everything bearing on the truth of P has been taken into account. In the most straightforward way, then, I might as well choose either one, as far as truth is concerned.

Background-relative permissivists are better situated. Knowing that I am in a permissive case, they may suggest, doesn't mean that my belief is arbitrary from the standpoint of truth, because my belief looks, antecedently, a lot more likely to be true, *from my background perspective*, than the belief of my disagreeing evidential peer. And, they may add, it is only against a background of priors or standards that we can even *begin* to make such a judgment. Though my evidential equal has a different view, from *my* perspective the standards or priors he has are worse, when it comes to the truth, (though not when it comes to rationality), than mine. (See Schoenfield 2014)

I think there is much to be said for this kind of reply. But it is easy to remain worried. Suppose you learned that the background standards or priors you accept were chosen for you in the following way: when you were young, a scientist came up with two backgrounds she knew would lead to directly opposite opinions regarding P. Then, she flipped a coin, and based on the outcome assigned the associated background to you, subsequently leading you to form the belief that P. Learning this would, plausibly,

justify you in losing confidence in your belief in P. It differs from a case where the scientist flips a coin and *directly* implants a corresponding belief in you, it seems, only by degree of elaboration. It is little comfort to notice that according to the standards the coinflip chose for you, the belief you reached looks pretty good. That looks like bootstrapping of a particularly flagrant variety. There is something arbitrary about the coinflip, which is insensitive to truth, that threatens to infect things all the way down. And since our adoption of backgrounds is, like the coinflip, not truth-sensitive, it looks arbitrary in the same way.

The non-evidential permissivist, finally, must claim that what makes the difference, sometimes, between rationally believing P and rationally believing its negation is some non-evidential factor. But evidence is precisely all that bears on the truth of a proposition. The non-evidential permissivist is no better off, *from the standpoint of truth* than an agent in an options-permissive case would be if they decided by a coinflip. They simply substitute for the coinflip some equally truth-irrelevant decisive factor.

I do not claim that the challenges facing the views I've raised above are unanswerable. But answering them is sure to take on controversial commitments which may turn a potential permissivist away. So it is worth seeing if other forms of permissivism with different, and at least for some, more palatable commitments can be found.

4.1 – Dynamic Permissivism

The three general forms of permissivism above are perhaps not exhaustive of the literature but they are representative. Each is associated with a distinctive way of responding to the evidence-pointing argument consistent with accepting a version of a *state-oriented* or *synchronic* principle – one judging us according to the states and combinations of attitudes we are in at particular times – that demands that at all times our beliefs be supported by our evidence. Such a principle is intuitively appealing, and implicit or explicit commitment to state-oriented norms of this sort are common (Broome 2007, Hedden 2015). So it is not surprising that the discussion of permissivism by both advocates and opponents has effectively

taken such principles for granted. If we reject the assumption that there is a state-oriented principle relating the beliefs of rational agents to their epistemic grounds, however, a novel approach to permissivism presents itself.

A *dynamic, process-oriented* norm takes as its jurisdiction not the question of which beliefs, intentions, or other mental states we may be in at any given time, as state-oriented norms do, but rather the operation of mental processes – temporally extended activities, like reasoning or consideration, through which we manage our cognitive economy.

These norms are not merely norms that mention states at multiple times, demanding, for instance that if we have some attitudes at a time t_1 , we must have some other attitudes at a later time t_2 . Though a norm like this is diachronic or intertemporal, it still takes our states as their fundamental object. Though this is one way of departing from the synchronic principle, the view I would like to explore goes further. Genuine process norms take *processes*, a very different kind of mental entity from a state like belief, as their object, telling us when and how a particular process operates in a rational agent, and how it interacts with the rest of our cognitive universe.

The dynamic view I have in mind, like non-evidential permissivism, denies premise (1) in the evidence-pointing argument. Also like that view, it recognizes the importance of a replacement principle connecting our beliefs to the considerations that support them. But instead of replacing it with another state-oriented norm policing our mental attitudes, it replaces it with a set of norms that are dynamic and process-oriented in the sense just described. While norms on attitudes are mostly limited to telling us which attitudes or combinations thereof we may hold, process norms come in at least two types – norms governing when a process is to be initiated, and norms governing how it is to operate once initiated. In particular, I want to focus on the process of *considering whether P is true*.

My proposed replacement for the *State-oriented Evidential Support* principle is the following:

Dynamic Evidential Support: It is rational to form or reaffirm a belief that P as the conclusion of a process of consideration-whether-P only if P is supported by the evidence one takes into account during consideration.

Two important questions must be answered before we can see the consequences of this kind of dynamic norm; First, when is a rational agent required to consider whether P? And second, what evidence must they take into account when considering? Not every possible combination of norms on consideration leads to an interesting form of permissivism. In particular, combining extremely demanding norms like “whenever an agent gets new evidence that has any bearing on whether P, they are rationally required to consider whether P” and “whenever an agent considers whether P, they are rationally required to consider their total evidence” will, with the principle above, lead agents to quickly recalibrate in line with other agents with the same evidence whenever they learn something new.

But I think those norms are overdemanding in at least two ways. First, it is not very plausible that any bit of relevant evidence is enough make it rationally required for us to consider a question – we learn many things all the time that have at least some small evidential relevance to countless beliefs, and we do not have time to consider all of them. Second, though perhaps more controversially, it is implausible we are required to be responsive to all of our evidence whenever we consider a question. Our total evidence, the full range of information we could, in principle, bring to bear on a question, is a very large set, and in considering what to believe, taking evidence into account is costly in both time and cognitive resources. Sometimes, limited agents like us must get by while taking into account only parts of our total evidence.

There is a tradition in formal epistemology that takes the norms of rationality to describe the

behavior of ideal agents not subject to the kinds of cognitive limitations faced by human beings⁵. Adherents to this ideal picture of rationality will not be moved by worries about overdemandingness. The debate over idealization in the study of rational norms is beyond the scope of this paper, so I will merely note that the view I am developing is not neutral regarding this point, and suggest only that the ordinary intuitions about rational behavior to which I appeal do seem to take constraints posed by our limitations seriously. And we will at least learn something about how theoretical commitments about the relation of rational norms to cognitive imperfection constrain the forms of permissivism available to us.⁶

If we take these thoughts on board, we are in a position to discover a path towards a robust kind of permissivism. In the following sections, I will make what I take to be some pretheoretically plausible claims about the norms on initiating consideration and on taking evidence into account, and use them to illustrate how, on a dynamic picture, rational agents with the same evidence may come to differ in their beliefs.

4.2 – Non-evidential Factors

My first suggestion is that *non-evidential* factors make a difference both to when an agent is required to consider whether P, and to how much of their evidence they are required to take into account. This can happen in a number of ways, but I will just mention two.

First, external practical pressures can preclude consideration or limit how much evidence the agent must consider. Suppose it is a typical Thursday, and an educational program on the radio I am listening to provides me with some evidence which I suspect has significant bearing on the thesis of a paper I am writing. This gives me sufficient reason to reconsider whether the thesis is true. Now imagine a case that is identical, except that I am receiving this new evidence through the radio at the same time that my apartment is in the process of being consumed by a fire. In this second case, I am required *not* to consider

⁵ See, for example, Christensen (2007) on “rational ideals”.

⁶ I owe an appreciation of this point to an anonymous referee.

whether my thesis is true. There are simply more pressing matters to attend to. It may be that in both cases, the new evidence I have acquired since the last time I considered the question is such that, were I to consider, I would change my mind on the topic. But because in one case I rationally consider and in another case I rationally do not, I will end up with different beliefs, despite my total evidence bearing on the question being the same (the fire, we are assuming, does not itself bear evidentially on the thesis of my paper).

Similarly, if I am at home considering something on a lazy day with little else to do, I may rationally take a great deal of time exploring my evidence when thinking about a question, and so am permitted or required to take a great deal of my evidence into account. On the other hand, if I have a doctor's appointment upcoming, I may only feasibly have time to take the most salient evidence into account. Suppose that deep down in my evidence, there is some *subtle clue* such that if I were to take it into account, it would change my mind on the topic. In the first case, where I have time to explore my evidence thoroughly, I rationally do take it into account, and in the other, I rationally do not. So I will end up with different beliefs despite my total evidence bearing on the question being the same (again, we may postulate that the question on which I am reflecting is evidentially independent of whether I have a doctor's appointment or a lazy day ahead).

Another way non-evidential factors can make a difference to rational consideration is by raising the stakes. Suppose an acquaintance tells me that she is pretty sure she saw Joe at a party last weekend, and I form the belief that he was there on that basis. Some time passes, and I've encountered subtle clues, in the form of passing comments by partygoers that would, if I considered the question and took those clues into account, be sufficient to undermine my confidence that Joe was at the party. But normally, we may imagine, I have no reason to reconsider such a question. Or, if I consider it, I have no reason to look deep enough into my evidence to notice the subtle clues. So I hold on to my belief. But now imagine that I am at a murder trial where the defense rests its case on the possibility that Joe, and not the defendant, committed

the murder on the very night I believe he was at the party. Now it is very important that I get the answer right. So I am rationally required to reconsider the question of whether Joe was at the party, and when I consider it, to take as much of my evidence into account as time reasonably allows, which may lead me to take into account the subtle clues I would otherwise have missed. Again, I might end up with different beliefs in the two cases, despite having the same evidence.

4.3 – Order Matters

My second suggestion is that the *order* in which we receive evidence can be relevant to the norms on considering and taking into account. One way this may happen is motivated by the following thought: receiving a large chunk of evidence relevant to a proposition often makes it rational for us to reconsider whether that proposition is true. But receiving a tiny bit of evidence relevant to a proposition does not. We just don't have time to consider all those questions when over ninety-nine times out of a hundred it will make no difference to what we believe. Now imagine two agents who begin with the same evidence and form a belief in P. One of the agents then receives a huge chunk of evidence which is, in total, sufficient to undermine his belief, and perhaps even form the opposite belief. The other agent slowly, once a month, over a period of twenty years, is fed tiny bits of evidence. Each bit in some barely significant way counts against P. After twenty years, this evidence in sum is equal to the chunk received by the first agent. It would, intuitively, be irrational for the first agent to fail to respond to the massive influx of evidence bearing on her belief. But it seems rational for the second agent, each month, not to respond to some tiny new evidence regarding P. Since each month she is rational in not responding to the negligible change in her evidence, she will still be rational after twenty years. Again, they end up evidentially symmetric but with different beliefs.

One might worry that I am taking advantage of some sort of Sorites-related difficulty here, but the similarity to Sorites cases is only superficial. We can assume it is perfectly clear at what point the balance of

the agent's total evidence begins to weigh against P. What's doing the work in this case is that whether one is required to *reconsider* whether P in response to some new evidence bearing on P, a precondition for taking into account the overall bearing of our evidence, is not a matter of how the agent's *total* evidence bears on P, but a matter of how significantly *the new bit of evidence* bears, or seems to bear, on P. For we should reconsider in response to some new evidence only if we think there's some good chance that that evidence *makes a difference*. Importantly, we do not generally keep track of all the times in the past some barely significant evidence bearing on a question came into our possession. And each decision to reconsider is taken independently. So if each new piece of evidence doesn't meet the threshold of seeming like it would make a difference, then we end up rationally never reconsidering even though bit by bit the justification provided by our total evidence has been largely siphoned away.

In another, similar case, suppose that there are three pieces of evidence, E1, E2, and E3, which bear on whether P. And suppose that individually, they are not significant enough that one is required to reconsider whether P upon getting that piece of evidence, but that any two of them together are significant enough that one must reconsider. Finally, assume that E1 and E2, together with my old evidence, justify believing P, but that E1 and E3, together with my old evidence, justify believing not-P. Now, if I get E1 *and* E2 as a single chunk of evidence, and then get E3 later, I will end up rationally believing P. If I get E1 and E3 as a chunk, however, and then get E2 later, I will end up rationally believing not-P, even though my total evidence, in the end, is the same in each case.

4.4 – Doing Better than You Need

The final way I'll mention that plausible norms on processes lead us towards a permissive picture comes from the thought that you're often not rationally *required* to consider a question or take some evidence into account, but that it is not a mark against you if you do. Suppose A and B are agents with the same evidence and beliefs regarding P, and neither is in a position where they are required to consider

them, but that if they were to reconsider, they would rationally change their mind. It is the weekend, and A doesn't have much better to do, so he considers whether P, and consequently changes his mind. B would rather spend her weekend playing tennis. So, they end up with different beliefs.

Alternatively, imagine A and B have the same evidence and beliefs, and are both considering whether P. Moreover, there is a subtle clue buried deep enough in their evidence that they are not required to notice it, but which, it turns out, is decisive in favor of not-P, while the less subtle evidence looks decisive in favor of P. Now suppose in the middle of her consideration, B notices a billboard that reminds her of the subtle clue. B, it seems, must take this clue into account and therefore form the belief that P, while A, permissibly not noticing the subtle clue, forms the belief that not-P. Once again, they find themselves with different ultimate beliefs.⁷

5.1 – Why Dynamic Permissivism?

We have seen a number of ways a dynamic view might allow evidentially symmetric agents to rationally diverge. The distinctive way in which the dynamic view works, I want to suggest, carries very different commitments from other permissive views and is especially well-placed to avoid some of their worries.

First, dynamic permissivism *does not entail that individual agents have options* about what to believe. Indeed, dynamic permissivism is consistent with a strong kind of uniqueness thesis:

Uniqueness of Evidential Support: Any set of evidence supports exactly one set of doxastic attitudes, for any agent.

⁷ One may be wondering, at this point, whether the sort of considerations just sketched make room for permissive *philosophical* disagreements. There is some reason to think that the extent to which they do is comparatively limited. This is because for people who are active participants in an academic dispute - it won't be the case that the agents differ because one of them rationally hasn't considered the issue, or hasn't considered it deeply. That comes with the job, so to speak. And those are the main ways in which agents can, on the view described, end up permissibly disagreeing on the same evidence.

Together with *Dynamic Evidential Worth Believing*, this uniqueness thesis means that there will never be more than one way for an agent to conclude their reasoning – the outcome is fully determined by the evidence they consider. I did appeal in the previous section to the thought that agents are allowed to be more conscientious than they need to be in considering a question or taking evidence into account. But the choice to think harder about a question or to dig deeper for evidence is a kind of *practical* permission – representing options about what to *do*. And claiming that agents with the same evidence cannot rationally differ in what they *do* would be a highly implausible position. So it avoids the objectionable features of options-permissivism.

Second, dynamic permissivism avoids commitment to *lower-level arbitrariness*. Starting point and standards permissivism were able to say that given an agent's non-evidential commitments, their beliefs were not arbitrary. But, while they might be able to make the unexciting claim that the agent's own standards or starting points aren't arbitrary from the standpoint of those standards or starting points themselves, there is a clear sense in which they are arbitrary from an objective or third-person point of view. Dynamic permissivism makes no such commitments.

Relatedly, dynamic permissivism has the *best response to the arbitrariness argument*. The key premise, we recall, is (B). Consider what it means, for the dynamic permissivist, to be in a permissive case. All that means is that there is some way of rationally considering and taking into account evidence, or not considering, or not taking into account evidence, by which someone may have come to a different opinion. But this may well happen because they *didn't consider* a question that we considered, or because they didn't take as much evidence into account. It seems clear that if that's the reason they have a difference of opinion, it gives me no reason to be skeptical of my own belief nor does it look, from a third-person point of view, as arbitrary from the standpoint of truth. It is much like learning that someone else who is less *informed* than I am rationally holds different views. It is only when I learn that my rational evidential peer holds a different view due to *more or equally* diligent consideration that should, plausibly, lead me to worry,

and therefore motivate me either to reconsider the question in depth or modify my confidence.⁸ So being in a permissive case is not enough to make one think one's belief is arbitrary from the point of view of the truth.

Next, dynamic permissivism is able to recognize a sense of *the supremacy of the evidence*. Although practical factors related to our human limitations can affect the attention we are able to give to a doxastic question, the dynamic view allows that in the end, what to believe is decided by a function of the evidence that I manage to consider alone. No non-evidential factors weigh *against* the evidence. If something about the thought that evidence should reign supreme when it comes to rationality of belief is right, this is an advantage over the non-evidential view.

Finally, *Uniqueness of Evidential Support* and *Dynamic Evidential Worth Believing* together represent, I believe, the best concession any form of permissivism is able to make to supporters of evidential uniqueness. Options-permissivism flouts any obvious kind of uniqueness thesis. Background-Relative Permissivism can accept that evidence uniquely justifies (though given the possibility of permissive standards, standard permissivism has some work to do), but only relative to an agent's commitments, which are rationally open. Cross-agentially, uniqueness disappears. The non-evidential permissivist can accept the uniqueness of evidential support, but must admit that the evidential support is not decisive. Dynamic permissivism, by contrast, can say that evidence points exactly one way, that it does so even from an objective, cross-agential point of view, and that it is in an important sense decisive for belief.

5.2 – Dimensions of Permission

Mileage may vary regarding how worrisome the reader may find the challenges for traditional

⁸ On the dynamic view, the disagreement of others might have bearing in at least two ways – first, on whether one reconsiders a question, and as potential evidence during the consideration itself. Plausibly, learning someone, particularly someone who has similar or better evidence, or considered a question as carefully as you did, disagrees with you is a reason to reconsider. And plausibly, the testimony of experts as well as of equals has some evidential force. But this leads us into contentious territory – the overall dynamic permissivist approach is strictly compatible with many different views about the evidential bearing of the beliefs of others.

permissivist views I've raised. I certainly don't take them to be conclusive. But I think there is good reason to get excited about dynamic permissivism even if one is absolutely convinced that there are genuine epistemic options, that there are many permissible standards or starting points, or that there are non-evidential reasons bearing on belief. I have presented these views as *alternative* permissivist models, but this is perhaps misleading, since they can, in fact, be combined. They are more like dimensions along which doxastic difference between evidential equals may be permitted. One may even hold all of them together. But it is important to at least make the distinction; lumping them all under the heading of "permissivism" obscures many important differences. Keeping them separate allows one to make subtle distinctions between types of permissive cases that may be important when mapping the contours of the full permissive landscape. Perhaps, for instance, *moderate* permissivism is true along one of the dimensions of permissivism, like the options dimension, and *extreme* permissivism is true along another, like the dynamic dimension.

Skeptics of permissivism should pay attention to these distinctions as well. As we've seen, different forms of permissivism are vulnerable to different sorts of objections, and manage to recapture uniqueness-friendly intuitions and principles to various degrees. Each may have its own peculiar vulnerabilities, but they differ in their structure in ways that make it unlikely that any blanket objection to permissivism as a whole will do the job.

Permissivism, I have argued, represents a diverse family of views with little in common except their name and their basic commitment to the existence of rational disagreement, and one paints them with the same brush at one's own peril. At the least, I hope that I have succeeded in producing a portrait of this family that accentuates those differences. After this, I hope that I have introduced the reader to a cousin they have never met before. And lastly, I hope I have motivated the thought that this cousin, odd though they may be, is worth getting to know a little better.

References

- Ballantyne, N. and Coffman, E. J. (2011). "Uniqueness, Evidence, and Rationality." *Philosopher's Imprint* 11(18): 1-13
- Christensen, David. (1994) "Conservatism in Epistemology," *Noûs* 28: 69 - 89.
- Christensen, David. (2007) "Does Murphy's Law Apply in Epistemology?"
Oxford Studies in Epistemology 2:3-31.
- Feldman, R. (2007). "Reasonable Religious Disagreement" in L. Antony ed. Philosophers Without God.
Oxford, UK: Oxford University Press: 197–214
- Hedden, Brian (2015). "Time-Slice Rationality." *Mind* 124 (494):449-491
- Harman, Gilbert (1986). Change in View: Principles of Reasoning. Cambridge, MA: MIT
- Kelly, T. (2005). "The Epistemic Significance of Disagreement" in T. Szabo Gendler and J.
Hawthorne eds. Oxford Studies in Epistemology Vol. 1: Oxford, UK: Oxford University Press:
167–196.
- Kelly, T. (2011). "Peer Disagreement and Higher Order Evidence" in A. Goldman and D. Whitcomb
eds. Social Epistemology. Oxford, UK: Oxford University Press: 183–220.
- Kelly, T. (2013). "Evidence Can Be Permissive". In Contemporary Debates in Epistemology, 2nd edition.
Matthias Steup and John Turri eds. Malden, MA: Blackwell Publishing
- Lycan, William. (1988). Judgement and Justification. New York: Cambridge.
- Meacham, Chris. (2013) "Impermissive Bayesianism." *Erkenntnis* 1-33: 1185-1217
- McGrath, Matthew. (2007) "Memory and Epistemic Conservatism." *Synthese* 157 (1):1 - 24
- Nelson, Mark. (2010) "We Have No Positive Epistemic Duties." *Mind* 119: 83-102
- Rosen, G. (2001). "Nominalism, Naturalism, Epistemic Relativism." *Philosophical Perspectives* 15:

69–91.

Schoenfield, Miriam. (2014) “Permission to Believe: Why Permissivism is True and What it Tells us About Irrelevant Influences on Belief.” *Nous* 48: 193-218

Schroeder, Mark. (2012) “Stakes, Withholding, and Pragmatic Encroachment on Knowledge”. *Philosophical Studies* 160: 265-285

Sklar, Lawrence 1975. “Methodological Conservatism.” *The Philosophical Review* 84, 3: 374-400.

White, R. (2005). “Epistemic Permissiveness.” *Philosophical Perspectives* 19 (1): 445–459