**Trusting yourself means trusting others: Why necessary self-trust does not generate agent-centered norms**

Word count: 7,800 (including footnotes, excluding references)

Abstract: Reasoning is impossible if we don't trust some of our own mental states in some sense. It is possible, however, even if we don't trust others' mental states of the same sort in the same way. Some philosophers have used this as the basis for an argument that epistemic norms are agent-centered: that they allow one to treat one's own mental states differently than those of another, even when one is aware of the other's mental states and has no reason to think that the other's mental states should be treated differently from one's own. Agent-centrism would have important implications for the debate on disagreement. However, arguments for agent-centrism made on the basis of the need for self-trust will run into one of three problems. They will license self-trust much less than is really appropriate, they will beg the question, or they will show that we should trust others in much the same way that we trust ourselves. Since the first two are fatal flaws, to the extent that these arguments work at all, they undermine the claim that epistemic norms are agent-centered. I show why this is, and discuss the potential ramifications for disagreement.

## 1. Introduction

Let's say that I know something about you: I know that it seems obvious to you that *p* is true. Now what – how should my beliefs change in light of this? Once I know of it, does the seeming to *you* that *p* have the same normative implications for me as the seeming to me that *p* would (whatever they are)? Similar questions can be asked about epistemically significant mental states other than seemings. Occasionally, these questions are easy to answer: sometimes

we know something about other agents that justifies our treating their mental states differently from our own. But that leaves open the question of whether, in the absence of such knowledge, agents should treat their own mental states differently from the mental states of other agents when forming or revising beliefs. Put another way, this question asks if there is something about the mental states that belong to an agent, other than the fact of that agent's awareness of those states, that allows or requires that agent to epistemically privilege those states.[1] Let's call epistemic norms that answer "Yes" to this question, at least sometimes, *agent-centered.*[2]

The question of whether epistemic norms are agent-centered is not only interesting in itself, it also has application to the ongoing debate on disagreement. Agent-centered epistemic norms lend themselves more readily to the claim that, in cases of apparently problematic disagreement, we can justifiably maintain our beliefs. Recently, two philosophers – David Enoch (2010) and Michael Huemer (2011) – have separately argued that epistemic norms are agent-centered, and both have applied their versions of agent-centered norms to the question of how to respond to disagreement. While they differ on which of our own mental states we can privilege, and under what conditions, the underlying arguments for their views are quite similar and represent an attractive line of thinking. This paper is about a flaw in that line of thinking.

The main idea behind the argument I will criticize is that reasoning or thought is

[1] The question of how we should treat others' mental states, and whether we should treat them differently from our own, really only arises when we know of others' mental states. For the rest of this paper, when I talk about how A should treat the mental states of B, I am assuming that A is aware of the relevant mental states of B (perhaps but not always because B has informed A of the relevant state).

[2] I'm adapting this term from Huemer, 2011.

impossible if we don't trust some of our own mental states – some of our own beliefs or opinions or intuitions or experiences – in some sense. However, this mental process *is* possible even if we don't trust *others'* mental states of this sort in the same way. This difference, it is claimed, licenses trusting our own mental states more than those of others even when we lack any particular reason to do so, which means that epistemic norms are to some extent agent-centered. Let's use the label *the argument from necessary Trust* to refer to arguments of this sort. I am capitalizing "Trust" here and hereafter to make salient that I'm not using the term in the ordinary sense. Any advocate of this argument will have something in mind that bears a resemblance to the ordinary notion of trust, but which will not be likely to map on to the ordinary notion exactly.

Different thinkers will articulate versions of the argument from necessary Trust that differ in their details. It will be helpful, though, to start by looking at a generalized version of the argument in premise-conclusion format:

1. For any agent A, there is some type of mental state M such that A must Trust some instances of M that belong to A to avoid some serious epistemic problem.
2. There are no mental states belonging to any other agent such that A must Trust them to avoid this serious epistemic problem.
3. For any two sets of mental states G1 and G2, if an agent must Trust some members of G1 to avoid this serious epistemic problem but need not Trust any of G2 to avoid the problem, then it is epistemically permissible for that agent to Trust members of G1 and to not Trust any member of G2.
4. Thus, for any agent A, it is epistemically permissible for A to Trust A's M states and not those of others.

To understand any specific version of this argument, we need to know what it means by "Trust,"

what mental states are to be Trusted, and what sort of serious epistemic problem Trust is needed in order to avoid. In the next two sections, I'll lay out Enoch and Huemer's formulations of the argument as paradigm instances of it. I will accept Enoch and Huemer's versions of premises 1 and 2 for the sake of argument. This is not because I believe all of what they say, but rather because I will go on to show that there is no version of premise 3 that will give them, or any proponent of the argument from necessary Trust, what they want. It turns out that any version of premise 3 will be false (even according to advocates of the argument), will beg the question, or will not support agent-centered epistemic norms.

### 2.1 Enoch's argument

Enoch argues that we should treat others as "truthometers" – instruments for determining the truth (as thermometers are instruments for determining the temperature).[3] According to Enoch, when you believe that *p* and I know that you do, this licenses my seeing *p* as true only when I can properly infer that *p* from your belief that *p*. The appropriateness of this inference depends on how reliable of a believer (how good of a truthometer) I take you to be. However, Enoch continues, I cannot generally treat myself as a truthometer. If I did, then when faced with the question of whether or not *p* was true, I would have to infer *p* from my own belief that *p*. This inference would require information about how reliable I am. If I were treating myself as a truthometer, determining how reliable I was would require inferring something about my reliability from my beliefs about my reliability. This inference would require further appeal to some belief about my reliability, and one can see how this generates a problematic regress.

What this shows, according to Enoch, is that I can't treat my own beliefs as evidence for

[3] All references to Enoch are from Enoch (2010) unless otherwise noted.

the truth of their contents; rather, when I believe that *p*, I should just treat *p* as true.[4] When I believe *p* and encounter a supposed peer who believes ~*p*, I can thus say, "Well, this person seems like a peer; however, they believe ~*p* even though *p*. So they can't be very reliable on this issue." In Enoch's own words:

> "[in a case where you believe *p* and Adam, a supposed peer, believes ~*p*]…given the ineliminability of the first-person perspective [the need to Trust one's own beliefs] and the (at least moderate) self-trust that comes with it, why on Earth should you *not* see Adam's belief *not-p* as reason to believe he is less reliable than you otherwise would take him to be? After all, when you believe *p*, you do not just entertain the thought *p* or wonder whether *p*. Rather, you really *believe p*, you take *p* to be true. And so you take Adam's belief in *not-p* to be a mistake. And, of course, each mistake someone makes (on the relevant topic) makes him somewhat less reliable (on the relevant topic) and makes you somewhat more justified in treating him as less reliable (on the relevant topic). Why should *this* mistake, then, be any different? … [D]oes taking *p* to be true, and using it as a premiss in an argument for demoting Adam from peerhood status, not simply amount to begging the question? No, it does not, or at least not in a problematic way." (Enoch, 2010, 979-980, footnotes omitted)

In terms of the argument from necessary Trust, the mental state that Enoch takes to be privileged is belief. The sort of Trust he has in mind is simply treating *p* as true when I believe *p*, without needing to make an inference to *p* from my belief that *p*. The serious epistemic

[4] Enoch doesn't intend this to be a completely unrestricted claim – for example, we might treat our past or future selves as truthometers, he says – but he does intend it to apply fairly generally to our present beliefs (see Enoch, 2010, 964-965).

problem to be thereby avoided is the sort of regress that leads to skepticism.

### 2.2 Huemer's argument

Huemer endorses Phenomenal Conservatism (PC), the view that one's own seeming that *p*, by itself, gives one *prima facie* justification to believe *p* in the absence of defeaters (Huemer, 2001). In his own words: "I believe, however, that a single principle can account for [the justification of] all foundational beliefs. The principle is… (PC) If it seems to *S* as if *P*, then *S* thereby has at least prima facie justification for believing that *P*." (Huemer, 2001, 99, footnotes omitted) He believes that PC must be true because all rational thought requires the truth of PC; if PC weren't true, then all premises used in reasoning would need to be reasoned to, which would lead to infinite regress. Further, Huemer claims, any argument against PC will be self-defeating, because the argument will always start from premises the arguer endorses because they seem true. However, he argues, rational thought and the avoidance of self-defeat do not require that others' seemings, by themselves, provide a basis for belief. We can be justified in treating others' seemings as evidence – and in fact in giving them as much credence as our own – but this requires some belief about their trustworthiness, whereas credence in our own seemings does not (Huemer, 2011). On Huemer's view, then, when our seemings conflict with those of another, and we lack positive evidence that the other's seemings are as reliable as our own, we are allowed to Trust our own seemings and not those of the other.

To put this in the terms I used in articulating the argument from necessary Trust, the mental state Huemer says that we must Trust are our seemings.[5] Seemings include intuitions and perceptual experiences, among other things. To Trust something, in Huemer's sense of the word, is to non-inferentially base beliefs on it. This is similar to the sort of Trust that Enoch has in

[5] Enoch explicitly denies this.

mind, since it involves simply accepting that some proposition is true without having to reason to the conclusion that it is true. The serious epistemic problem that Trust is necessary to avoid is either the impossibility of rational thought, or self-defeat.

**3. The problem for the argument from necessary Trust**

The argument from necessary Trust starts from the point that we must Trust some of our own mental states to avoid epistemic disaster, but need not trust any of the mental states of others to do so. It goes from there to the conclusion that it is permissible to not Trust others' mental states. The problem for the argument is in how it gets from this starting point to this conclusion: I will argue that any way of doing so will be faced with a trilemma.

In the generalized version of the argument I gave above, the premise that makes this step is premise 3. As I have it, premise 3 says:

> 3. For any two sets of mental states G1 and G2, if an agent must Trust some members of G1 to avoid this serious epistemic problem but need not Trust any of G2 to avoid the problem, then it is epistemically permissible for that agent to Trust members of G1 and to not Trust any member of G2.

Premise 3 doesn't clearly give us agent-centered norms, as it is compatible with my treating my own and others' mental states equally. As used in the argument from necessary Trust, G1 refers to the set of my beliefs or seemings, and G2 refers to the set consisting of the beliefs/seemings of others. However, if I need to Trust some members of the set of my beliefs/seemings to avoid regress or skepticism, then I also need to Trust some members of the set of *all* beliefs/seemings. This means that the antecedent of premise 3 would be true if G1 referred to the set of all beliefs and G2 to the set of all guesses (or hopes or dreams, etc). Thus, according to premise 3, it is not only permissible for me to Trust my beliefs and not those of others, it is also permissible for me

to Trust all beliefs, including those of others (and likewise for seemings).

Why is this a problem? After all, premise 3 apparently does allow one to *not* Trust others' mental states. To see the issue, imagine someone who only Trusted every other rational intuition of theirs, or only fourth-order beliefs about reliability, without any reason to not Trust their other beliefs or intuitions. They would be doing something wrong. Generally speaking, when all members of some set of things have epistemic relevance, it's impermissible to treat only an arbitrary subset as if they mattered. This is reflected, for example, in the notion that rational belief must be rational in light not just of *some* of the available evidence, but in light of one's *total* evidence (see e.g. Kelly, 2008); if some set of my evidence is relevant to a question, it's irrational/impermissible for me to take only some of it into consideration without good reason. Premise 3 on its face says that I may Trust either just my own mental states of a certain kind, or all mental states of that kind. It does not give me reason to choose between Trusting one or the other. If premise 3 licenses arbitrarily Trusting just my own mental states, then it is false. If it does not, then it is incomplete because it fails to articulate why I should Trust my own mental states over those of others, or it does not generate agent-centered norms because it requires me to Trust all mental states of the relevant kind (since to Trust any smaller subset would be arbitrary).[6] In any case, it is overly broad in what it allows.

Note that nothing I said above is specific to Enoch or Huemer's version of the argument.

[6] Note that changing premise 3 to say that we are required to Trust G1 and not G2 makes things worse. It would generate conflicting requirements. On Enoch's view, the antecedent of premise 3 is true for me if G1 is my beliefs and G2 is your beliefs, which would mean that I'm required to Trust my beliefs and none of yours. The antecedent is also true if G1 is all beliefs and G2 is all guesses, which would mean I'm also required to Trust *all* beliefs.

What Huemer, Enoch, or anyone who makes an argument from necessary Trust, need is a different premise 3. However, what we'll see is any version of premise 3 runs into one horn of a trilemma: like the version of premise 3 we have just discussed, it will be too broad, failing to explain why we should differentially Trust our own versus others' mental states (not supporting agent-centered norms), or allowing/requiring Trust in mental states that clearly should not be Trusted; it will be too narrow, having the result that we need not Trust mental states that we clearly should (or that advocates of the argument believe we should); or it will beg the question against those who don't advocate agent-centered epistemic norms.

### 3.1 Refining the argument from necessary Trust

Premise 3 of the argument from necessary Trust has to explain why we may Trust our own mental states and not those of others by pointing to differences between them. There are three sorts of ways to try to fix premise 3. One might continue to distinguish between our own and others' mental states based just on differences in the need for Trust, but do so more carefully than the premise 3 above does. One might try to distinguish between our own and others' mental states partly by appealing to inherent differences between them. Or one might try to make this distinction on the basis of differences in our relationship to mental states. I'll consider each of these below. For each of the proposed versions of premise 3 I give, changes will need to be made to premise 1 and 2 of the argument I give above to maintain validity; I'll assume for the sake of argument that these changes are unproblematic and won't discuss them.

#### 3.1.1 Individuating more carefully

The problem with premise 3 as stated is that it fails to really capture the idea behind the argument from necessary Trust. This idea is that we *must* Trust our own mental states but need not Trust others', and that this is an important difference between our own and others' mental

states. Premise 3 above, however, does not restrict Trust to just those mental states that must be Trusted. This might suggest the following version of premise 3 is more in line with the spirit of the argument:

> 3a. For any two sets of mental states G1 and G2, if an agent must Trust all members of G1 to avoid a serious epistemic problem but need not trust all of G2 to avoid that problem, then it is epistemically permissible for that agent to Trust members of G1 and to not Trust any member of G2.

3a does not work, however, because it is too restrictive. For whatever type of mental states that we plausibly should Trust, we only ever must Trust *some* instances of these states to avoid epistemic disaster. That means that we could divide these types of mental states into sub-classes, Trust one sub-class and not the other, and still avoid the problem that the argument is concerned with. In fact, it's quite likely that, for people like us, there is no set of mental states all of which must be Trusted. This would mean that the antecedent of premise 3a would never be satisfied, and 3a would not show that it is permissible to Trust only our own beliefs or seemings. To illustrate these problems, let's look at 3a in light of Enoch and Huemer's arguments.

For Enoch, "Trust" means "treat as true without inference." Enoch points out that inferring *p* from my belief that *p* requires consideration of my reliability, and if I always needed to infer my reliability from beliefs about my reliability, I'd get a problematic regress. I could avoid this regress, however, if I only Trusted my beliefs about the reliability of my beliefs. That is, whenever I had a belief about the reliability of my other beliefs, I simply treated it as true, while, for all my other beliefs, I inferred that they were true on the basis of the reliability I saw them as having. In fact, I could avoid *infinite* regress by just Trusting some of my beliefs about reliability: for example, Trusting only beliefs about the reliability of my beliefs about the

reliability of my first-order beliefs. So, neither the set consisting of my beliefs, nor of just my beliefs about reliability, satisfies the antecedent of 3a. What's more, our beliefs are often quite coherent. Because of this, specific beliefs of ours can typically be derived from some of our other beliefs. So, if there were a set of beliefs all of which plausibly must be Trusted, we could likely replace any member of the set with some other beliefs coherent with it. This would still allow us to avoid infinite regress, as we could derive or infer the replaced belief after Trusting the replacement members. This strongly suggests that there is no set of beliefs all of which *must* be Trusted to avoid infinite regress.

We see much the same thing with premise 3a when applied to Huemer's argument. Rational thought is possible, and self-defeat avoided, even if we only base beliefs on some of our seemings. Trust in intuitions alone, for example, without Trusting perceptual seemings, would make rational thought possible. Further, not *all* intuitions need to be Trusted: rational thought would be possible if we only based beliefs on every other (or every third, etc.) intuition.[7] External world skepticism, another worry that motivates PC, would be avoided if we based belief on only every other perceptual experience (or every third or every fifth, etc). Finally, our

[7] We can imagine people for whom full rationality is impossible without Trust in all intuitions: people who only have (true) intuitions about axioms such that each must be believed for fully rational thought, and only have each of them once. This wouldn't help Huemer's argument. The epistemic disaster his version of the argument from necessary Trust is about is the impossibility of any rational thought, and such a person could be fairly rational without Trusting all of their intuitions. In any case, we aren't people like that, since we have intuitions on multiple occasions, and also have "redundant" intuitions (so that we could derive the axioms from just some of them).

seemings are very often coherent. So, for any candidate for the set of seemings all of which need to be Trusted, we can likely replace one member with some other seemings coherent with it and avoid epistemic disaster. We should expect that there is no set of seemings all of whose members must be Trusted.

Even if turned out that there is some set of beliefs or seemings whose members all must be Trusted, these sets will be much smaller than either Enoch or Huemer would like. Enoch appeals to the argument from necessary Trust in order to argue that, for a very broad range of propositions *p*, when I disagree with an apparent peer about *p*, I am allowed to Trust my belief that *p* and take the fact of disagreement as evidence that the person is not a peer.[8] Huemer, similarly, believes that seemings, including perceptual seemings, are generally to be Trusted. Further, Enoch and Huemer are right to want to extend Trust fairly generally to our mental states of the relevant kind. If either of their views is to be plausible at all, it must say that very wide ranges of beliefs or seemings are Trust-worthy. In ordinary circumstances, it would be quite odd for a thinker to only treat a very restricted subset of their beliefs as true, or to only base beliefs on a very restricted subset of perceptual seemings. So premise 3a is subject to the second horn of my trilemma. It is too restrictive: at best it does not capture the range of mental states that Enoch and Huemer think should be Trusted, and at worst (and most likely) its antecedent is unsatisfiable.[9]

[8] Perhaps not for any *p*, but certainly he wants to include a philosophically interesting propositions. Since we could avoid regress by not Trusting these and just Trusting beliefs about reliability, these would not be Trust-worthy on 3a.

[9] One might point out that 3a as stated seems to allow one to Trust more broadly, since it does not say that we are required to not Trust mental states that don't satisfy its antecedent. If it does

Premise 3a might be on to something, however. An advocate of the argument from necessary Trust will want to restrict Trust to members of sets of mental states smaller than those that include the mental states of everyone. Perhaps we can do so simply by appealing to the idea that the size of a set of mental states makes a difference to whether its members should be Trusted. Normally size is not of much epistemic relevance, but I do think it makes sense to appeal to it here: we might not want to be overly Trusting, and so perhaps we should restrict Trust to just where it is needed. This suggests that premise 3 could be modified as follows:

> 3b: For any agent A and token mental state M, A is permitted/required to Trust M if M is a member of the smallest set G of mental states such that, if A were to Trust no member of G, then A would be subject to some serious epistemic problem; otherwise A is permitted/required to not Trust M.

I have stated 3b with "permitted/required" to allow the premise to be adapted to avoid different problems (which I'll discuss below). 3b seems to avoid the first horn of the trilemma by not being overly inclusive; because the set consisting of A's mental states of the relevant type is smaller than the set consisting of all mental states of this type, 3b avoids the problem I raised for premise 3. It also apparently avoids the second horn by not being as restrictive as 3a: it allows Trust in sets *some* of whose members must be Trusted. Even so, we will see that 3b runs into the same sorts of problems as 3 and 3a.

The set consisting of all of my beliefs (or seemings) is not the smallest set some of whose

---

allow Trust more broadly, however, then it falls into another horn of the trilemma: it allows Trust of others' mental states as well, and gives no reason not to do so. This would make Trust only in one's own mental states apparently arbitrary.

members I must Trust to avoid disaster.[10] Take the set consisting of all allegedly Trust-worthy token mental states of the relevant type; now subtract one member. The result will be a set such that, if I Trust none of its members, epistemic disaster follows (even if I were to Trust the one member subtracted). For example, if Huemer is right about the role intuitions play in reasoning, then if I were to Trust only a single one of my intuitions, rational thought would be largely impossible. And, if Enoch is right at all about Trust, my reasoning would almost always result in infinite regress were I to Trust only a single belief of mine. The choice of members that I may subtract to get this result might not be totally unrestricted, but certainly most of the members of the starting set are possible candidates. This means that there will be a vast number of sets of beliefs or seemings, some of whose members must be Trusted to avoid disaster, which are smaller than the set of all of our beliefs or seemings. We can repeat this subtraction process over and over (obviously not indefinitely), each time obtaining sets of beliefs or seemings some of which must be Trusted. This means that the smallest set of mental states some of which *must* be Trusted is much smaller than the set of mental states all of whom Enoch or Huemer think *may* be Trusted. It also suggests (although I can't prove this here) there may be no unique set of mental states that is the smallest set some of which must be Trusted.[11]

To see why this makes 3b unacceptable, we must consider different possible understandings of the normative claims 3b makes. There are two normative claims in 3b: one

[10] Enoch does not claim that all of my beliefs are to be Trusted; when I say "all beliefs," treat this as referring to "all beliefs that Enoch thinks are to be Trusted."

[11] This problem should generalize to other plausible accounts of Trust: for whatever set of token mental states is claimed to be Trust-worthy, we should expect there to be some much smaller sub-set of this set such that not Trusting any of its members will still be epistemically disastrous.

about Trusting members of set G (the smallest set such that an agent must Trust some of its members) and one about Trusting non-members. Each of these can be understood as stating a requirement or as making a claim about what is permissible. If either makes a claim about what we are required to do, and there is no unique smallest set G, then we will have serious problems. Since few if any mental states will belong to every one of the smallest Gs, if we were required to not Trust non-members of G, then for most token mental states there would be a G that generated a requirement to not Trust it. For the same reason, if we were required to Trust members of G and permitted to not Trust non-members, we'd be both required to Trust and permitted to not Trust most or all of our own mental states. 3b would also be problematic if both of its normative claims were permissibility claims. If they were, 3b would say that we may Trust members of G and we may Trust non-members of G. In this version of 3b, the distinction articulated makes no normative difference, and thus does no work in the argument from necessary Trust. Yet the whole point of the argument is that there is a normatively significant distinction between our own and others' mental states. So, if there is no unique G, 3b is too broad or too narrow. It either does not differentiate between our own and others' mental states, demands us to not Trust most of our own mental states, or makes consistent normative claims about (at best) a very small set of our own mental states.

Let's put that to the side and assume that there is some unique G. Even if there is, neither Enoch nor Huemer will accept that the second normative claim in 3b states a requirement. If it did, it would forbid Trusting non-members of G, and we've just seen that G will be smaller than the set of mental states who members Enoch and Huemer think may be permissibly Trusted. To elaborate on a point from above, one need not believe in agent-centered norms to agree with them that this would be too restrictive. If 3b restricts which beliefs may be Trusted, then there

would be large classes of our beliefs that *p* such that, *whenever* we wondered whether or not *p*, we would be obligated to reason as such: "I believe that *p*, and I'm reliable, so *p*." Infinite regress to the side, this sort of behavior looks nothing like ordinary thought, or like what ordinary thought should look like. Likewise, anyone with foundationalist leanings should be inclined to say that we may Trust more of our perceptual or intuitive seemings than would be included in the smallest G.

So for 3b to be plausible, it must state a permissibility claim about Trust in non-members of G: it must say that we are allowed to Trust or not Trust beliefs/seemings that don't belong to G. For reasons discussed above, however, we shouldn't be allowed to Trust arbitrarily. So, given 3b, a good reasoner should either Trust all of the mental states of the relevant type that don't belong to G (in the absence of reasons not to), or none. If the first were appropriate, then the argument from necessary Trust would fail to show that epistemic norms are agent-centered. If the second were called for, agents would be Trusting fewer mental states than advocates of the argument from necessary Trust think they should. 3b is either too narrow or too broad.

So far we've considered some ways of differentiating our own and others' mental states based just on the need for Trust. I have not considered every possibility, but I have covered an appropriate range of them, from the very broad premise 3 to the very narrow 3a, to the "just right" 3b. None of these on their own avoid the trilemma. We should conclude that this method of fixing premise 3 on its own won't work. Let's briefly turn to another alternative: adding to premise 3 an appeal to inherent differences in the mental states discussed.

### 3.1.2 Inherent differences

One of the problems with the versions of premise 3 discussed above is that they apparently allow people to Trust arbitrary subsets of mental states. The obvious response to

make is to give a version of premise 3 that only talks about non-arbitrary sets: "real" types of mental states. That's the idea behind premise 3c:

> 3c: If G1 and G2 are different natural kinds of mental states, and Trusting some of G1 is necessary to avoid some serious epistemic problem for agent A, but Trusting any of G2 is not, then it is permissible for A to Trust all of G1 and none of G2.

This would allow all of an agent's beliefs or seemings to be Trusted if any of them must be, thus avoiding the second horn of the trilemma. However, it doesn't get us agent-centered norms. It falls prey to the first horn of the trilemma, since it does not differentiate between our own and others' beliefs or seemings. This is because, while beliefs or seemings might be natural kinds, there is no natural kind that includes my beliefs but not yours, or my seemings but not yours.

I'm not going to discuss appeals to *inherent* differences between our own mental states and those of others any further. The advocate of agent-centered epistemic norms appeals to differences in necessary Trust to motivate agent-centrism exactly because there *aren't* inherent differences between our own and others' mental states. Appeal to inherent differences thus will not rescue premise 3 or the argument from necessary Trust. Instead, we should be considering differences in our relationship to our mental states versus those of others.

### 3.1.3 Relational differences

Let's start with the simplest possible appeal to the difference in our relationship to our own mental states versus those of others:

> 3d: All of an agent A's mental states (of a certain kind) but none of any other agent's, should be Trusted by A.

This would give Huemer and Enoch what they want, but it also clearly begs the question, which is the third horn of the trilemma faced by the argument from necessary Trust. What's more, both

Enoch and Huemer believe that it begs the question; both deny that being mine is epistemically relevant.[12] Rather, they believe that there is something that is true about my mental states, in virtue of their being mine, that that explains why I should treat them differently from those of others. That being mine is generally epistemically irrelevant is quite plausible, since there are a great number of mental states that I should not Trust, even though they belong to me: guesses, hopes, wishes, etc. This reasoning shows that we can eliminate any version of premise 3 that appeals only to relationships we have to too wide a range of our own mental states – e.g. that they are directly accessible, or transparent, to us – since guesses, hopes, and so forth will also stand in these relationships to us.

Are there special relationships we have to our own our mental states, but not others', that *are* epistemically relevant? One suggested by Huemer (personal communication) is incorporated into the following version of premise 3:

> 3e: If it is impossible for an agent A to Trust mental states in set G2, but it is possible for A to Trust mental states in set G1, then members of G1 but not G2 must be Trusted by A.

One might wonder why Enoch or Huemer would appeal to this, since at first glance it seems possible to Trust others' mental states: Fred could say that it seems to him that *p*, or that he believes that *p*, and because of this George could come to see *p* as true without reasoning about *p* at all. On second glance, however, it's not so clear that this would be Trust. Whenever George

[12] "I agree that it is unreasonable to give your own view extra weight simply because it is yours… That it is yours seems epistemically irrelevant." (Enoch, 2010, 986) Huemer never says this as explicitly, but it's clear that he accepts it from the arguments he gives for agent-centered norms.

comes to believe what Fred believes or what seems true to Fred, it's quite plausible, or perhaps unavoidable, that there is some mental state of the relevant kind that "comes between" Fred's belief/seeming and George's belief. Before Fred's belief can generate George's, George might have to believe that Fred believes that *p*. Similarly, it might have to seem that George that it seems to Fred that *p* (or that Fred has said that it seems to him that *p*) before Fred's seeming can be a basis of George's belief. In either case, one might argue that this intervening mental state prevents George from Trusting Fred's mental states. If that were right, then 3e would differentiate our own from others' mental states.

This is too quick, however; 3e is still stuck on one of the horns of the trilemma. To see how, let's consider two interpretations of the claim that we cannot Trust others' mental states. On one interpretation, Trust is impossible in interpersonal cases because of facts about causation: George can't Trust Fred's mental states because there is a mental state of George's that comes between (causally speaking) his beliefs and Fred's mental states. On the other interpretation, the right causal connections can exist, but *appropriate* Trust is impossible. I'll start by discussing the causal interpretation.

The causal interpretation of 3e is too restrictive. The psychology of the mind is complex, and we should be open to the possibility that all, or a great deal, of the mental processes that are paradigmatic cases of self-Trust involve complex or multi-step causal chains. That is, we should not be very surprised if what looks like non-inferential basing of beliefs on our own seemings, or treating of our own beliefs as true, typically or always causally involved mental states that came between the Trusted mental state and result of Trust. On the causal interpretation of 3e, if this unsurprising psychological fact obtains, disastrously few mental states can be Trusted, likely resulting in skepticism. Even if things aren't so bad, self-Trust is still too likely to involve

complex causal chains for Enoch or Huemer to stomach the causal interpretation of 3e. Consider the sorts of cases Enoch is most interested in: noting that a supposed peer disagrees with you about *p*. Such experiences will typically bring to mind second order beliefs, among them the belief that you believe that *p*, and these might very well play some causal role in the resultant treating of *p* as true. For Huemer, seemings are automatic and beliefs are not, which means that, when a belief of mine is based on a seeming of mine, there will often be some mental event that comes between the seeming and the belief (see Huemer 2001, 97). I doubt that Enoch, Huemer, or any advocate of the argument from necessary Trust would want to rest their argument on tenuous empirical grounds, or commit to the claim that Trust is impossible in worlds (among which may be our own) in which self "Trust" involved multi-step causal chains. The causal interpretation of 3e is too narrow.

The rejection of the causal interpretation should not be surprising. Theories of Trust are motivated partly by worries about regress. But the concern is about regress of justificatory, not causal, reasons. It makes sense that 3e must be interpreted in this light. 3e might say that, when some mental state comes between Fred's seeming and George's belief, it is possible but inappropriate for George to Trust Fred; alternately, Trust might be an inherently normative notion, and impossible where inappropriate. But premise 3e is supposed to explain why we should treat our own mental states differently from those of others, and simply saying that it's permissible to non-inferentially base beliefs on the one but not the other is begging the very question we are trying to answer. So premise 3e will either fall to the second horn of the trilemma – restricting Trust too much – or it will fall to the third by begging the question.[13]

[13] The ideas here obviously intersect with debates on testimony and about the basing relation. It may be that the right account of testimony or of basing will show that we can't non-inferentially

I have not discussed every possible difference in our relation to our own mental states versus those of others. However, the points I have made should apply generally: these differences will be irrelevant to Trust, because we stand in these special relations to mental states that we should not Trust; they will be too restrictive, as we won't stand in these relations to all of our own mental states that should be Trusted; or they will beg the question, because they will only seem epistemically relevant to those who already accept that epistemic norms are agent-centered. Thus they will either face some horn of my trilemma.

To summarize: if we say that we should Trust only mental states *all* of which must be Trusted, then we get the result that few or no mental states should be Trusted (the second horn of the trilemma). If we say we should Trust mental states *some* of which must be Trusted, then we either still get the result that too few mental states should be Trusted, or the result that others' mental states should be trusted (the first horn). If we try to appeal to inherent properties of mental states to explain which should be Trusted, we have to say that others' mental states should be trusted (again, the first horn). If we try to appeal to the relationship between us and our mental states, the result is either overly restrictive or begs the question (the third horn of the trilemma). As none of the individual differences we've discussed save premise 3, we might try find a version of premise 3 that combines multiple of the factors we've discussed. Since appealing to inherent differences does nothing, we'd have to build our premise partly based on

---

base our beliefs on others' mental states. If so, the argument for that account will give us agent-centered epistemic norms (assuming ought implies can), since it will show that others' mental states can't give me the same sorts of reasons that my mental states do. However, this does not count in favor of the argument for necessary Trust, as it would be the argument about basing or testimony that did all the work in establishing agent-centrism.

an appeal to differences in our relationships to our own versus others' mental states. However, since these are either epistemically irrelevant or beg the question, a premise that contains them would inherit these flaws.

**4. Conclusion and application to disagreement**

While I've focused on Huemer and Enoch's arguments, I've tried to make my points sufficiently general that they will apply to any plausible version of the argument from necessary Trust. The different versions of premise 3 that I have discussed, covering as they do tweaks not only to the appeal to the need for trust, but also to inherent and relational properties, should cover the range of ways of fixing premise 3 that one might consider. The main assumption needed to generalize my arguments is that, for any notion of Trust necessary to avoid epistemic disasters like skepticism, plausible theories will have to extend this notion to plenty of token mental states that strictly speaking don't *have* to be Trusted. Because others' mental states are so similar to our own, once we do so we will either have to extend Trust to others' mental states or endorse a question begging theory. This means that, if any version of the argument from necessary Trust holds any water, it is best seen as showing that we should Trust others' mental states in much the same way we Trust our own. This doesn't prove that we should reject agent-centered norms. Instead, it shows that one initially plausible argument for them instead is really an argument against them.[14]

[14] Richard Foley (2005) and Linda Zagzebski (2007) separately have argued for conclusions superficially similar to mine, although neither intends to (nor does) address the argument from necessary Trust. Both argue, roughly, that the rationality of self-trust implies or leads to the rationality of trust in others. Both arguments show that we should sometimes *inferentially* base our beliefs on the mental states of others, which is consistent with both Enoch and Huemer's

If the argument from necessary Trust does show that we should Trust others' mental states, what does this tell us about disagreement? It would mean that, in circumstances in which we should Trust our own and another's mental states, how we should respond to disagreement depends on what we should do in the face of our own incoherence. To illustrate, consider first Enoch's approach. When I disagree with a peer about whether *p* or ~*p*, since I should Trust the peer's belief as well as my own, I should act as if *p* and also act as if ~*p*. This will also happen when I myself believe both *p* and ~*p*. Whatever response is appropriate in this latter situation should also be appropriate in the case of peer disagreement. I suspect, although I have no space to explore the idea, that this will lead to an account of disagreement that looks a lot like the Equal Weight view (Elga, 2007). If it does, then there is an argument for the Equal Weight view that sidesteps the objections to the view that Enoch gives in the paper I've been discussing, since the argument is in essence based on Enoch's argument *against* the Equal Weight view.

Similarly, if one adopts Huemer's view, then to figure out what to do when one's seemings differ from a peer's, one has to ask what one should do when one intuits contradictory things. Should we suspend judgment on the two intuitions until we have further evidence on the topic – does each act as a defeater to the justification provided by the other? Or perhaps should we believe whichever is stronger? I won't attempt to address these questions here.

One's own beliefs or experiences typically play a different role in one's epistemic life than the beliefs or experiences of others. It's hard to imagine not relying on one, but easy to imagine not relying on the other. This makes natural the argument that there is an epistemically relevant difference between them, a difference which licenses treating them differently in our reasoning and belief formation. We should resist the temptation to make this argument. We

views.

should resist it not because our beliefs or intuitions are not special, but rather because, if they are, *everyone's* are.

**Acknowledgements**

[removed for blind review]

**References**

Elga, A. (2007) Reflection and Disagreement. Nous, 41,478–502.

Enoch, D. (2010) Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement. Mind,119, 953-997.

Foley, R. (2005) Universal Intellectual Trust. Episteme, 2, 5-11.

Huemer, M. (2001) Skepticism and the Veil of Perception. Palgrave McMillan.

Huemer, M. (2011) Epistemological Egoism and Agent-Centered Norms. In Dougherty, T. (ed.) Evidentialism and its Discontents, Oxford.

Kelly, T. (2008) Evidence: Fundamental concepts and the phenomenal conception. Philosophy Compass, 3, 933-955.

Zagzebski, L. (2007) Ethical and Epistemic Egoism and the Ideal of Autonomy. Episteme, 4, 252-263.