

May 2014

From Self-trust to Other-trust: the Role of Reasons and the Theoretical Insignificance of Behavioral Inconsistency

Alex Papulis

University of Wisconsin-Milwaukee

Follow this and additional works at: <https://dc.uwm.edu/etd>



Part of the [Epistemology Commons](#)

Recommended Citation

Papulis, Alex, "From Self-trust to Other-trust: the Role of Reasons and the Theoretical Insignificance of Behavioral Inconsistency" (2014). *Theses and Dissertations*. 419.
<https://dc.uwm.edu/etd/419>

This Thesis is brought to you for free and open access by UWM Digital Commons. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of UWM Digital Commons. For more information, please contact open-access@uwm.edu.

FROM SELF-TRUST TO OTHER-TRUST: THE ROLE OF REASONS
AND THE THEORETICAL INSIGNIFICANCE OF BEHAVIORAL INCONSISTENCY

by

Alex Papulis

A Thesis Submitted in
Partial Fulfillment of the
Requirements for the Degree of

Master of Arts
in Philosophy

at

The University of Wisconsin-Milwaukee

May 2014

ABSTRACT
FROM SELF-TRUST TO OTHER-TRUST: THE ROLE OF REASONS
AND THE THEORETICAL INSIGNIFICANCE OF BEHAVIORAL INCONSISTENCY

by

Alex Papulis

The University of Wisconsin-Milwaukee, 2014
Under the Supervision of Professor Edward Hinchman

There is intuitive pull to the idea that if a subject knows or believes that others are just like her, and if she trusts herself as an epistemic agent, then she should also trust others as epistemic agents. It seems that she would be inconsistent to trust herself, and then when faced with an identical being in a situation or environment identical to her own, not extend trust to that other being as well. In this paper I argue that charges of inconsistency of the sort above can only apply if one's self-trust is on the basis of reasons. I consider why the intuitive consistency claim sounds right, address some possible concerns about my conception of inconsistency, and then go on to suggest that behavioral inconsistency is theoretically insignificant – it needn't play any significant role in a theory of trust intended to show that other-trust requirements are not being met. Theories wishing to show that there are other-trust requirements should give reasons for trust. I conclude by considering an objection involving self-trust policies.

Thanks to my advisor Edward Hinchman, and to all the participants of the UWM Spring 2013 seminar on self-trust and higher-order evidence. Thanks also to Paul Blaschko for leading a reading group of Zagzebski (2012) in the summer of 2013 and, additionally, for giving me feedback on a draft of a paper that served as the basis of this longer one. Thanks also to all who contributed thoughts during the UWM Fall 2013 philosophy grad student writing workshop.

TABLE OF CONTENTS

Introduction	1
I. Inconsistency	4
II. Trust without anything taken as an epistemic reason and charges of inconsistency	7
III. The intuitive claim	10
IV. Other conceptions of inconsistency	13
V. The theoretical insignificance of inconsistency	18
VI. Prelude to a conclusion	22
VII. A concluding objection	23
Works Cited	29

There is intuitive pull to the idea that if I know (or believe) that others are more or less like me, and if I trust myself as an epistemic agent in certain situations, then I should also trust others as epistemic agents in those same situations. There is even more pull to the idea that if I know (or believe) that others are *exactly* like me and in exactly the same kind of situation or environment I'm in, then if I trust myself as an epistemic agent, I should also trust those others. It seems that I would be inconsistent to trust myself, and then when faced with an identical being in a situation or environment identical to my own, not extend trust to that other being as well. If I have a set of faculties that is shared by others, and if these faculties are responsible for the formation of my beliefs, it seems that something is wrong if I rely on my particular faculties but not on those same faculties in another subject. If there's no difference in the faculties, it seems that I should have to rely on them equally in whatever subject or subjects they may be found, all things being equal.

An intuitive claim:

If a subject trusts herself as an epistemic agent, then she must trust as epistemic agents those whom she believes to have epistemic faculties identical to her own (in identical environments)¹, on pain of inconsistency.

This is a plausible principle, but as I will argue in this paper, it is not available on just any account of trust in one's own faculties. That is, the principle doesn't hold apart from certain further conditions, which may or may not be met in individual cases of self-trust. So, if a general account of what makes self-trust an acceptable epistemic practice fails to acknowledge these further conditions², this intuitive principle will not

1 I will tend to leave off the condition of "identical environment" in the discussion that follows. My argument doesn't hinge on anything here, and for full effect, reference to identical (or similar) environment may be inserted wherever there is reference to identical (or similar) faculties.

2 I take it that Foley (2001), (2005) and Zagzebski (2012) give this sort of account. That's not to say that their accounts can't be amended or clarified so that inconsistency charges can be made.

hold without qualification on that account. I will argue that accounts conceiving of self-trust as something that subjects do (or just as well may do) without taking certain things to be epistemic reasons to self-trust are unable to ground charges of inconsistency against just anyone who would withhold trust to others.

An analysis of inconsistency will be critical to showing this. Inconsistency in multiple instances of one's behavior, I will suggest, is not merely a matter of difference in behavior in otherwise similar situations, but rather the failure to treat something as a reason, when one already takes that thing to be a reason and has acted/believed on the basis of it. Given this understanding of inconsistency, inconsistency charges with regards to other-trust can only be made against those who, in trusting themselves, take something to be a reason for their trust, and yet fail to treat that thing as a reason when it would require them to trust others.

The structure of this paper is as follows. In the first section I examine behavioral inconsistency charges in general, arguing that a subject's reasoning, i.e. taking things to be reasons, is the basis of such charges and that without a subject taking things to be reasons for belief or action there cannot be any problematic sort of inconsistency when different behaviors are exhibited.

While it's *necessary*, if there are to be inconsistency charges, that a subject who withholds trust to others take herself to have reasons for her self-trust, in the second section I show that taking oneself to have reasons for self-trust is not a *sufficient* condition for inconsistency charges. This is because some things which might be taken as reasons to trust oneself aren't applicable as reasons to trust others, and therefore one can't *fail* to treat those things as reasons when not extending trust to others. These first two sections, then, serve to demonstrate that the possibility of behavioral inconsistency

requires that not just any reasons, but reasons of a certain sort be involved in self-trust.

At this point in the paper, the goal will shift away from showing the importance of reasoning for making inconsistency charges and towards supporting the claim that inconsistency is theoretically insignificant, in the sense that theorizers who wish to show that other-trust requirements aren't being met by subjects have little need of the concept of behavioral inconsistency in their theories. In the third section, which lays some groundwork for this, I suggest that the intuitive claim this paper began with sounds right due to a hidden assumption about the reasons why subjects' faculties produce true beliefs, namely that there's some law-like connection between the sort of faculties one has and the tendency for a subject to form true beliefs. In the following section I address objections to the account of inconsistency briefly described above and formulate a more general condition of behavioral inconsistency. This supports the point of the next section: inconsistency is theoretically insignificant, because there's always something other than inconsistency that an inconsistent subject is doing and which is wrong. This conclusion, along with the discussion of the intuitive claim with which this paper begins, is leveraged into a suggestion for a different approach to trust accounts which aim at requiring other-trust: depending on what a subject takes to be his or her reasons, withholding trust may not be appropriate; so, accounts should focus on giving reasons, which we can accept, for why we should trust (however much we should). In other words, showing that trust is *actually* required will mean actually giving subjects reasons to trust. I conclude by considering an objection involving self-trust *policies*.

I. Inconsistency

Before demonstrating that charges of inconsistency don't always apply (without further conditions holding) when a subject withholds trust to others whom she knows or believes to have faculties like her own, I will discuss how inconsistency does and does not work. The sort of inconsistency we're particularly interested in is between instances of trusting and withholding trust. So, considered more generally, what we're interested in is inconsistency in behavior (what I'll sometimes call "inconsistency between instances of behavior" or "behavioral inconsistency"), as opposed to inconsistency between behavior and belief or inconsistency in sets of beliefs. By inconsistency in behavior, then, we mean something that at least entails (but is not merely identical to) "in situation Y, sometimes doing X and sometimes not doing X." On its own, this is not a complete description of the inconsistency we're concerned with, which is something that would somehow be problematic, something that one can be *charged* with. This means we must be concerned with more than just behavioral variety and uniformity in otherwise identical circumstances.

To illustrate this point, imagine a person getting out of bed one morning, choosing to put her right foot forward before her left, for no (contrastive) reason. The next morning, nothing relevant having changed in her situation, she chooses to put her left foot forward, for no (contrastive) reason. Is our subject *inconsistent* because one day she gets out of bed with her right, and the next with her left, even though she is confronted with exactly the same situation on both days? It seems clear that she is not doing anything wrong; there's nothing we can charge her with. Faced with identical situations, she chose different courses of action. Her behavior exhibited variety versus uniformity, and yet she was not inconsistent.

I suggest that inconsistency depends largely on the role that reasons play (or don't play) in a subject's behavior. It's beyond the project of this paper to give an account of reasons, but I will explain in broad terms how I am using the term and give a few examples to illustrate what I mean, though what follows is admittedly incomplete. By reasons I mean things on the basis of which we should act or believe. For example, if I should eat given that I am hungry, then my hunger is a reason for me to eat. By epistemic reasons, I mean those things on the basis of which we should form specific beliefs. The ringing of my alarm, then, would be a reason to believe that it's morning, my great track-record in counting to ten a reason (perhaps) to think my counting-to-ten faculties are reliable. I will also use "epistemic reason" in a derivative sense. Something may be an epistemic reason to trust oneself, when this thing is an epistemic reason (in the original sense above) to believe that one's faculties are accurate. In that case, my great track-record counting to ten might be an epistemic reason to trust myself when I'm counting small numbers of objects. By practical reasons, I mean those things on the basis of which we should act. For example, if I am thirsty, my thirst is (perhaps) a reason to get a drink of water. I will rely heavily on the idea that a subject can *take something to be a reason*, and I will call acting or believing on the basis of something taken to be a reason (either mistakenly or not) *treating that thing as a reason*.

To illustrate the relevance of reasons to charges of inconsistency, let's slightly modify the situation we considered above: now, our subject believes her right leg is a little stronger than her left and that she is likely to stumble if she steps away from the bed with her left first. Given this additional element, things may be different, depending on whether or not reasoning is involved (at least explicitly, and perhaps implicitly, as well) in her choice. If our subject chooses on Day 1 to get out of bed with her right leg first

because she believes it's stronger, etc., then she would be inconsistent if the next day, nothing having changed, she chose to get out of bed with her left. She would be inconsistent on Day 2, failing to act according to what she takes to be a (conclusive) reason, namely the comparative strength of her right leg (and, perhaps, her beliefs/desires regarding the consequences of stepping out of bed with either leg). Alternatively, we might say that she is failing to *treat* the comparative strength of her right leg as the reason she in fact takes it to be. She takes her right leg's comparative strength to be a reason to step with that leg first (that's why she stepped the way she did on Day 1). But treating that as a reason on Day 2 would mean stepping first again with her right leg, and so in failing to do so, she is failing to treat her right leg's comparative strength as the reason she takes it to be.

So, it appears that inconsistency is very tightly connected to what a subject takes to be reasons and how she treats those things in making decisions. In fact, inconsistency seems to be a bad thing because it amounts to failing to treat as a reason something that one thinks to be a reason. A fuller exposition of inconsistency would consist of more than examples, but it's clear that inconsistency in behavior is more than difference in behavior, and it's plausible that presence or absence of reasons is relevant to the applicability of inconsistency charges. Difference in behavior isn't problematic in itself; ignoring or failing to do what one thinks is required of oneself in some sense or other *is*. Later, in Section IV, I'll consider some objections to this specific conception of inconsistency.

II. Trust without anything taken as an epistemic reason and charges of inconsistency

I'll now argue that accounts which conceive of self-trust as something a subject does (or might as well do) without taking anything as an epistemic reason for this trust cannot categorically ground charges of inconsistency for withholding trust to others whom one knows or believes to have similar epistemic faculties. In other words, such accounts cannot ground charges of inconsistency *simply in virtue of* self-trust and the withholding of trust to others who are believed or known to be similar.

If a self-trust account conceives of self-trust as something that occurs (or might as well occur) in the absence of anything being taken to be an epistemic reason, then on that account a subject trusts herself not on the basis of anything taken as an *epistemic* reason to believe she is (at least generally) a true-belief-former.³ For example, a subject not trusting on the basis of an epistemic reason wouldn't be trusting herself on the basis of what she believes to be a good track record in forming true beliefs.

With self-trust thus restricted, we can make a distinction between two different sub-categories: self-trust on the basis of a practical reason, and self-trust on the basis of no reason at all. Accordingly, we can determine whether or not inconsistency arises, and with what modality, in cases of each sub-category of self-trust. Beginning with self-trust on the basis of practical reasons, the key question will be whether or not practical reasons that are (taken to be) present for oneself are (or would be) also applicable as reasons when considering others as potential objects of trust. If they are, then self-trust on the basis of practical reasons will be able to ground charges of inconsistency: withholding trust to others will amount to failing to treat as a reason something that one takes to be a

³ Here I gloss over what would make an epistemic reason concerning one's faculties, i.e. a reason to believe something about one's faculties, a reason to trust oneself. While it seems, for example, that (the belief in) the infallibility of one's faculties would be a reason to trust those faculties, I don't address this issue in this paper.

reason and has treated as such. If, however, practical reasons for trusting oneself aren't *necessarily* applicable when considering others as potential objects of trust, then withholding trust won't *necessarily* amount to inconsistency – a subject won't necessarily be *failing* to treat as a reason what she previously took to be and treated as a reason.

In fact, practical reasons for trusting oneself are not necessarily reasons for trusting others. This is clear if we imagine an example. If a subject's reason for trusting herself is to avoid skepticism, there are ways to avoid skepticism without trusting others, in which case the subject's principle of skepticism-avoidance can't be one undergirding a charge of inconsistency. Once she treats skepticism-avoidance as a reason to trust herself, skepticism can be avoided, making skepticism-avoidance no longer applicable as a reason to trust others⁴, and then the subject needn't be *failing*, in withholding trust to others, to treat skepticism-avoidance as the reason she takes it to be and previously treated it as.

So, not all practical reasons are sufficient for inconsistency, but that fact doesn't prevent some practical reasons from being sufficient. A couple of examples will serve to demonstrate that some practical reasons can be sufficient for grounding inconsistency charges. If, say, a subject's (conclusive) reason for trusting herself is to contribute to her goal of maximizing trust as a general phenomenon in the world (odd, yes, but not incoherent), then it would be inconsistent for her to trust herself for this reason and yet upon encountering another possible object of trust to fail to trust (barring other relevant considerations). Similarly, if a subject trusted herself because she wanted to advance in her career, it's likely that withholding trust to others would be inconsistent with her

4 It is possible that a subject *could* continue to take the principle of skepticism-avoidance as a reason (and not be confused; she might believe, say, that there could come a time when having beliefs formed on the basis of trust in others would be necessary for avoiding skepticism), but it certainly isn't something that one *necessarily* will have to continue taking as a reason.

self-trust. If career advancement is a (conclusive) reason for self-trust, then in most cases it will be a reason to trust others (at least to some extent), and so withholding trust to others would constitute a failure on the subject's part to treat something (career advancement) as the reason she takes it to be and previously treated it to be.

The second way, we said above, that a subject may trust herself (given the absence of anything being taken to be an epistemic reason) is in the absence of anything at all being taken to be a reason, epistemic or practical. In this case, given the preceding argument about the necessity of reasons, it's immediately clear that there can be no charge of inconsistency: if one trusts herself for no reason, she can't fail to take into account something she both takes to be a reason and previously treated to be a reason when faced with other potential objects of trust.

It turns out, then, if we're not considering the possibility of trust on the basis of something taken to be an epistemic reason, that charges of inconsistency can be made only contingently, depending on the practical reasons, if any, that one has for trusting oneself. This means that sometimes, and perhaps even often, charges of inconsistency *will* be appropriate. However, since the only reasons we're allowing a subject (for now) to take herself to have for self-trust are practical reasons, inconsistency will be dependent on the presence of practical reasons, and that doesn't seem to do justice to the intuitive claim at the beginning of this paper.

The intuitive claim we began with doesn't reference reasons of any sort, and intuitively it seems like we shouldn't have to reference reasons. The next section will concern why the intuitive claim sounds right, and this explanation will be useful when focus of this paper turns to the questions of inconsistency's theoretical importance and of how theorists interested in requiring self- and/or other-trust might go about doing so.

After I argue that inconsistency is relatively insignificant in terms of showing that other-trust requirements aren't being met, I'll suggest that the best way to require trust is to try and give reasons for thinking that human faculties can be regularly accurate. The results of the examining the intuitive claim we began with will support this suggestion.

III. The intuitive claim

The intuitive claim this paper began with might be pressed as follows: if what is responsible for my belief is just like what is responsible for your belief, I can't think that my belief is somehow better off than yours. And that point doesn't depend, an objector might add, on whether or not self-trust is on the basis of any reasons. This sounds reasonable, and I think there is a common and, in this case, hidden notion that is motivating this reasoning. This notion is that there is some feature or something about what is responsible for my belief that explains *why* it produces the beliefs it does. In other words, there's a reason why my beliefs are the way they are, and this involves the properties and characteristics of the faculties from which they issue. Such a notion makes the intuitive reasoning above sound reasonable: if my faculties produce beliefs in the way they do *because of* the way they are, then if your faculties are the same as mine (and in the same situation), they will produce beliefs just like mine do, and if that's the case, I can't think that your faculties, *ceteris paribus*, produce different beliefs than mine do.

At this point, though, all that we've said about inconsistency still blocks charges of inconsistency for not extending trust: if I don't rely on my faculties for any reason, then I'm not inconsistent by withholding trust to others. What bears remarking on here is that this is the case *even if* I think that my faculties would produce the same beliefs that

yours are producing if I were in your situation (given, still, that I don't trust my faculties for some reason). This is because, there being no reason involved, my self-trust amounts to taking the *actual* output of *my* faculties as accurate, but it doesn't require me to take any stand on my faculties' *potential* output. This is certainly a counter-intuitive idea and we would think someone odd for behaving this way, but it's not incoherent. If I don't believe there's a reason why my faculties get it right (such as some law-like connection between my faculties' properties and their accuracy), nor therefore that there's a reason why my faculties in some non-actual situation *would* get it right, then I'm not obviously committed to taking any stand, one way or the other, on my faculties' potential output or the output of faculties with the same characteristics as mine. So, while it seems strange to say that a subject can acknowledge that another's beliefs are the same as those that her own faculties would yield, and yet withhold trust, this needn't be problematic if both 1) my trust is not on the basis of a relevant reason, and 2) I don't hold a relevant belief about the *nature or cause of the accuracy* of my faculties in producing beliefs.

This discussion, suggests, then, that somewhere behind the intuitive claim we began with is the assumption that there *is* a law-like connection between the character of a subject's faculties and their accuracy. Of course, if the assumption isn't the assumption or belief of the subject to whom the inconsistency claim is being applied, but is rather an assumption of a third party (another subject, perhaps, who is making the charge of inconsistency), then the intuitive claim will not in fact apply. On the other hand, if the assumption is the self-trusting subject's, then there may be something wrong with withholding trust.

One of the uses of this section is that it helps us see why the intuitive claim at the beginning of this paper sounds right. Additionally, the results of this section will lend

support to the suggestion I will make at the end of the paper: the way to require self- and other-trust isn't by relying on inconsistency charges, but by offering reasons to think human faculties are regularly accurate, such reasons that we in fact will accept while all the while admitting that we are using these very faculties in finding, offering, and accepting reasons. If a subject believes that there is a law-like connection between the operation of human faculties in certain situations and the accuracy of the beliefs that result, and if as a result she is doing something wrong if she doesn't extend trust to bearers of human faculties, then perhaps the way to show that other-trust is required is to *make it required* by finding, offering, and accepting reasons to rely on human faculties.

At this point it will be useful to look at some objections centered on my conception of inconsistency. Addressing objections will lend support to my position, and it will also bring us close to an important intermediate point that supports my suggestion that the best way to show that other-trust is required involves finding and giving reasons to rely on human faculties. This intermediate point is that inconsistency between instances of behavior is relatively theoretically insignificant. Looking at objections to my conception of inconsistency will require refining the conception, and this in turn will be useful in later highlighting the fact that behavioral inconsistency's wrongness derives from something other than the inconsistency itself. And because what it's derivative from, as we'll see more starkly in this section, are certain facts about a subject's treatment (or non-treatment) of things as reasons, I'll argue in Section V that behavioral inconsistency is theoretically insignificant.

IV. Other conceptions of inconsistency

If the account of inconsistency I've suggested above fails to describe certain cases that we think are cases of behavioral inconsistency, that's a strike against the account and, by extension, a strike against what I've said about the relation between self-trust, other-trust, and inconsistency. One might wonder if there aren't cases of inconsistency where, contra my account, a subject *isn't* failing to treat as a reason what she takes to be reason. I'll now consider two kinds of cases which are at least initially plausible as cases of inconsistency. The two kinds of cases each involve instances of dissimilar behavior in otherwise similar situations, but in neither kind of case is the inconsistent behavior an instance of a subject not treating as a reason what she takes to be a reason.

One kind of case: instead of a subject in the first instance of behavior treating as a reason what she takes to be a reason and then in the second instance failing to treat that same thing as the reason she takes it to be, in the *first* instance of behavior she *fails* to treat as a reason what she takes to be a reason, and then in the second instance of behavior she treats as a reason that thing she takes to be a reason (but previously failed to treat as such). For example, we might think a person inconsistent who, after eating a bucket of fried chicken, goes out of her way to prepare something healthy for her dinner guests, because, say, healthy food is the best option when one can afford it. In the first instance, the subject fails to treat as a reason what she takes to be a reason, and in the second she does treat as a reason what she takes to be a reason.

When we look at the first instance of behavior in this case, we can't say that the subject is inconsistent, as there isn't any other instance of behavior yet for it to be inconsistent with (again, the sort of inconsistency we've been concerned with is inconsistency between different instances of behavior). If we had seen the subject of our

example finishing off a drumstick before beginning to plan for her guests the next day, we wouldn't think her behavior of inconsistent with something else she'd done. As for the second instance of behavior, it isn't a case of the subject failing to treat as a reason that which she in fact takes to be a reason, because as we've described this kind of case, the second act is when the agent *does* treat as a reason that thing that she takes to be a reason. So, on my account of inconsistency, the subject is inconsistent neither in the first instance of behavior, nor in the second instance. And, yet, such cases are plausibly cases of inconsistency.

There are multiple ways I can respond. First, it's not clear that this actually is a case of inconsistency. When a subject fails to treat as a reason what she takes to be a reason, but then later *does* treat it as the reason she takes it to be, *in that second instance of behavior* it doesn't seem like she's doing something wrong. And if being inconsistent entails doing something wrong, then the agent isn't being inconsistent.

Second, I can concede with little cost that what consideration of these sorts of cases shows is that a charge of inconsistency needn't be necessarily applicable at the time of an instance of behavior. I am inclined to think this isn't correct, but, if it is correct, it hardly affects anything I've said. Rather, it would make these sorts of case straightforwardly fit my account of inconsistency, with a slight modification: an inconsistent subject fails to treat as a reason what she takes to be a reason, and *at some point* (though, not necessarily a point in the past) she *does* treat that thing as the reason she takes it to be. So, the subject in such a case would on my account be inconsistent in the first instance of behavior, as it is an instance of failing to treat as a reason something that she takes to be a reason (and which at a later time she *does* treat as the reason she takes it to be).

Similarly, I can concede with little cost that what consideration of these sorts of cases shows is that inconsistency needn't be understood as attributable to a single action. Instead, perhaps inconsistency can be something that a subject engages in *across* or *in* multiple instances of behavior. This isn't simply to restate the sort of inconsistency that we're interested in – inconsistency between instances of behavior. Nor is it merely to say that one engages in inconsistency (partially) *in virtue of* engaging in multiple instances of behavior, though this would be correct on this concession. Rather, it's to say that inconsistency is predicable of a set of instances of behavior a subject engages in. If nothing I've said has depended on inconsistency being attributable to a single action, then this concession doesn't hurt my position, and it costs only a slight modification: a set of instances of behavior is inconsistent when an agent in one instance treats as a reason what she takes to be a reason and in another instance (we'll say either before or after), she fails to treat that reason as the reason she takes it to be. It's still the case that no reasons means no inconsistency, and the heart of the inconsistency is still failure to treat as a reason what one takes to be a reason.

Another sort of case: a subject 1) treats as a reason what she takes to be a reason (let's call that thing X), 2) at another point, either before or after, she doesn't treat X as a reason, but 3) neither does she then take X to be a reason; however, 4) the change in what she takes to be a reason isn't justified. This sort of case would look similar to cases of inconsistency such as have so far been considered, but what's crucially missing is the subject taking some one thing in *both* instances of behavior to be a reason. What might such a case look like? Generally, we might think of such cases as ones in which an subject exhibits instability with respect to what she takes to be reasons. This is what the fourth condition above is supposed to capture: what one takes to be a reason changes

unjustifiedly. As an example, then, we might imagine a subject who one summer takes (and treats) high transatlantic ticket prices as a reason to stay home over the summer, and then the following summer, without thinking about and/or coming to learn something relevant, no longer takes high ticket prices to be a reason and takes a trip to London. This sort of case, though, isn't limited to instances in which the subject unjustifiedly "loses" a reason. A subject might, instead, unjustifiedly "gain" a reason. She could, one summer, not take (and not treat) high transatlantic ticket prices as a reason to stay home, instead buying tickets and taking a trip to London, but then the next summer refuse to fly anywhere, now taking high ticket prices as a reason to stay home, even though she hasn't thought about and/or come to learn anything relevant about what she had taken to be her reasons.⁵

It's at least initially plausible that such a subject is behaving inconsistently. Yet, as she is not failing to treat as a reason what she takes to be a reason, my account of inconsistency doesn't describe this sort of case as one of inconsistency. There's a way of modifying my account to make it more accommodating of alternative and/or broader conceptions of inconsistency.⁶ The basis of this is a common element that the second kind of case considered above shares with cases of inconsistency as I've conceived of it, and that's that (at some point) the subject is *failing* to appropriately treat (or not to treat) something as a reason. The disjunction here is necessary, since the second sort of inconsistency case just considered allows that one might unjustifiedly *begin to take*

5 A similar sort of case, which I won't go into in this paper, is one in which the change in what the subject takes her reasons to be isn't unjustified, but in which what the subject initially takes and/or doesn't take to be her reasons is unjustified. I suspect that these sorts of cases will be harder to characterize as cases of inconsistency.

6 I'm inclined to think that these cases aren't really cases of inconsistency between instances of behavior, what's wrong really being the instability/unjustifiedness of what the subject takes to be her reasons. I'll set this aside, though.

something as a reason. If we allow that a subject is failing to appropriately treat (or not treat) something as a reason partly in virtue of her unjustifiedly taking (or not taking) that thing to be a reason⁷, then in the second sort of inconsistency considered above, the subject is failing to *appropriately* treat (or not treat) something as a reason.

In cases of inconsistency as I've conceived of it, the subject isn't treating as a reason *what she takes to be a reason*. In the most recent class of cases, the subject either 1) isn't treating as a reason something that, unjustifiedly, she no longer takes to be a reason, or 2) is treating as a reason something which, unjustifiedly, she has come to take to be a reason. In any of these cases, the subject isn't appropriately treating things that she can take to be reasons. However, there's no problem if we don't think that there's a real common element between the two sorts of inconsistency. It's enough for later purposes that in the first sort of case, a subject isn't treating as a reason what she takes to be a reason, and in the latter sort, a subject changes what she takes or doesn't take to be reasons unjustifiedly. We can allow a disjunctive necessary condition for inconsistency, though I'll continue as if there is a true common element in all cases of inconsistency.

So, taking into account the objections and responses above: if a subject is inconsistent (with respect to multiple instances of behavior), then in one of those instances she is not appropriately treating something which she can take to be a reason. This is a necessary condition for inconsistency, though (it would seem) not a sufficient one. That it's a *relatively* contentless condition makes it flexible: if inconsistency (or,

⁷ I realize that this would be a substantive claim. Furthermore, it raises the question of whether *any* treatment of unjustified reasons is appropriate. It would seem that one is doing something wrong in not treating an unjustified reason as a reason, but here I'm also suggesting that one is doing something wrong in treating an unjustified reason as a reason. As I note later in the text, though, it's not ultimately important that I be able to characterize all cases of inconsistency in one way, so long as for any given case of inconsistency a simpler "bad" element can be identified. So, I can allow that the treatment of reasons in these "instability of reasons" cases of inconsistency isn't bad, but the unjustified change is.

perhaps, the *sort* of inconsistency we're interested in) involves in a particular way reasons, or the treatment of reasons, or changes in reasons, then that can be used to give content to the "appropriateness" part of the condition. So, what's inappropriate in an inconsistent subject's treatment of reasons may depend on *whether* she takes something to be a reason, for example, or it may depend on *why* she takes something to be a reason, or it may depend on *when* she takes something to be a reason.

Responding to possible objections has left us with a pared down condition for inconsistency, and this condition centers on a subject's treatment of what she takes (or at least could take) to be her reasons. In the next section, I will argue that inconsistency between instances of behavior is theoretically insignificant, in the sense that inconsistency needn't have a role in trust theories. The pared down condition developed here will be useful, as it is not a sufficient condition. This means that it cannot be constitutive of inconsistency, and because this non-constitutive condition of inconsistency is itself a failure for the subject, whenever there is inconsistency there will be something other than inconsistency which is itself bad.

V. The theoretical insignificance of inconsistency

If we want a theory of why subjects must trust others, there's reason to suspect that inconsistency needn't play an important role in such a theory, and this stems from the analysis of inconsistency initially developed earlier in the paper and further developed in the last section. Given the analysis of inconsistency developed here, in cases of inconsistency there will always be a failure to appropriately treat (or not treat) something as a reason. What exactly makes treatment or non-treatment appropriate will depend on the precise sort of behavioral inconsistency that we're considering, but whatever sort that

is, it's still the case that there's a failure. Crucially, this failure isn't a sufficient condition for inconsistency between instances of behavior, as the condition we've developed doesn't require that there be multiple instances of behavior. Because it's not a sufficient condition, the condition itself must be something other than inconsistency; inconsistency doesn't simply amount to failing to treat (or not treat) appropriately something that can be taken as a reason.

For each of the sorts of behavioral inconsistency considered above, we can describe an analogous kind of case where one fails to appropriately treat (or not treat) something as a reason, but in which there is *no other instance of behavior* in otherwise similar circumstances to make it inconsistency. And, yet, this analogous sort of case would still be *bad*; the subject would still be doing something wrong.

Let's rework some previous examples to see this. Joan slightly injured her right leg, and the very first time she has to get out bed since injuring her leg, even though she takes her slight injury as a (conclusive) reason to get out with her left leg first, she decides to get out of bed with her right leg first instead. The next day her legs feel equally strong, and so she never is faced with the same choice. Joel wants to go to London, and while initially he thinks nothing of the high ticket prices, for no good reason his view of the ticket prices changes, he takes the high ticket prices as a reason to skip the trip, and he decides to stay home for the summer. Or: Joel doesn't want to go to London. Apart from not wanting to go to London, he takes high ticket prices as a reason not to go London. Then, for no good reason, his view of high ticket prices changes, and when he later desires to go to London, he buys his tickets.⁸ In all three of these examples, we have

8 It may not be possible to "subtract" an instance of behavior and reformulate cases of behavioral inconsistency (of the "instability of reasons" sort) where a subject in the first instance takes something to be a reason and that (putative) reason is such that it is necessarily treated as such. An example of the sort of case I have in mind: a subject takes high ticket prices as a reason not to buy tickets, and so

one instance of behavior, to which a second instance *could* be added to give us inconsistency between instances of behavior, but which are on their own cases of subjects doing something wrong.

This result suggests that inconsistency is theoretically insignificant, because whenever there's behavioral inconsistency, the subject is doing something non-identical with inconsistency (but which is related to the inconsistency) and which is itself something wrong. If there's always something else that one is doing wrong when engaging in behavioral inconsistency, then maybe the theory needn't focus on inconsistency. If the theory can work with simply the failure to appropriately treat (or not treat) certain things as reasons, then inconsistency is theoretically insignificant.

The sort of inconsistency and the sort of theory this paper is concerned with have to do with self-trust and other-trust. There's good reason to think that theories concerned with trust, and particularly concerned with showing that subjects are required to extend trust to others, could work just as well by focusing simply on an inconsistent subject's failure to appropriately treat (or not treat) certain things as reasons. We can see this by considering what cases of inconsistency would look like and how a theory that wasn't concerned with inconsistency would deal with them.

The sort of cases of behavioral inconsistency involving self- and other-trust will generally look as follows: a subject trusts herself, and when confronted with another subject who she believes to have similar faculties and to be in a similar environment, she does not extend trust. We know that reasons have to be involved somewhere, and there

simply in virtue of *not buying tickets* she is engaging in what I have called an "instance of behavior". As a consequence, a later change in behavior accompanying a later unjustified change in what one takes to be a reason will necessarily be a *second* instance of behavior. This isn't a problem for me, since it will still be the case that part of what constitutes the inconsistency will, *considered on its own*, be something bad for the subject.

are three options: 1) the subject takes herself to have a reason to trust herself, and also takes herself to have reason to trust the other subject; 2) the subject takes herself to have reason to trust herself, but no longer takes herself to have this reason (i.e. “loses the reason”) when faced with the second subject; and 3) the subject doesn’t take herself to have reason to trust herself, but when faced with the second subject does take herself to have reason to not trust. There are only these three options, as the self-truster either takes herself to have a reason for self-trust (in which case she can retain or lose this), or she doesn’t take herself to have a reason (in which case she either “gains” a reason or continues to take herself to have no reason to trust). If the agent never takes herself to have a reason, inconsistency can’t be at play, leaving the three possibilities.

We can see in each of these three possibilities that there is something the subject is doing wrong, which is more basic than the inconsistency. In case 1, the subject takes herself to have a reason to trust the second subject, yet she withholds trust; in case 2, the subject unjustifiedly ceases to take herself to have a reason to trust when faced with the second subject (and doesn’t treat this thing as a reason); and in case 3 the subject unjustifiedly begins to take herself to have a reason to withhold trust when faced with the second subject (and does treat this thing as a reason in withholding trust).

All this strongly suggests that a theory to require other-trust needn’t reference inconsistency. A subject is doing something wrong in failing to appropriately treat (or not treat) something as a reason, and inconsistency is dependent on this failing. Focusing on this failing is sufficient for generating charges against withholders or self-trust, so long as subjects take themselves to have reasons.

VI. Prelude to a conclusion

The preceding discussion has shown three things: 1) without a subject taking herself to have certain reasons, her self-trusting doesn't imply a consistency requirement to extend trust to others, because 2) reasons are central to inconsistency charges; and 3) since reasons are central to inconsistency charges, inconsistency in fact isn't all that important for theories that would require other-trust. Inconsistency shouldn't be relied on to require other-trust, and in this final section I'd like to suggest a way of refocusing efforts to require other-trust.

This paper began with an intuitive claim, and if we can preserve it (though perhaps not its *seemingly* straightforward logic), it's worth seeing how we might do that. Inconsistency isn't all that important, but reasons appear to be, or at least what we take to be our reasons. I'd like to suggest a reason-centered approach for developing accounts which require other-trust: if one wants extending trust to be required of epistemic agents, then she should develop an account which gives reasons for thinking that epistemic faculties tend to produce true beliefs. As we saw in our discussion of why the intuitive claim at the beginning of this paper sounds right, the idea that the character of one's faculties and environment is responsible for the accuracy of beliefs produced can make the intuitive claim work. So, if an account of why human faculties get it right can tie their character and the character of the environment in which they operate to the accuracy of the beliefs they produce, then those who accept the account will be required to trust others (and themselves).

Something to note, though, is that if a subject accepts an account of why her faculties are largely accurate, she'll be doing something wrong if she fails to trust, *regardless* of whether or not she takes herself to have a reason to trust. It would seem

that, if she believes that human faculties are such as to produce true beliefs in certain environments, withholding trust would be wrong on that account.⁹ An account of a theorizer seeking to require other-trust, then, ought to describe the law-like connections between various human faculties and the various situations those with human faculties find themselves in on the one hand, and the various levels of accuracy of beliefs produced in these different situations on the other.

If what I've said up till now is correct, then the way to require other-trust is to give reasons for it. There's no easy way for requirements to be spontaneously generated, so to speak; what a subject believes does seem to be relevant. This means that, if someone wants there to be requirements, the way forward is to affect what subjects believe – and that's a matter of giving them reasons to think that humans are generally well equipped in their environments for getting things right.

VII. A concluding objection

I'd like to close by considering a two-part objection. The first part claims that my conception of theoretical (in)significance needs to be developed to allow for an ignored sort of theoretical significance, and the second part of the objection claims that, given this development, a plausible view about self-trust *policies* provides an example of how inconsistency may be theoretically significant.

One may concede that inconsistency is theoretically insignificant insofar as something other than inconsistency can always be charged of an inconsistent subject, while also thinking that it is theoretically significant in a compatible way. A

9 I'm taking it for granted that someone who believes that he is in a situation in which his beliefs are (probably) true (or would be true if he didn't suspend belief), but who withholds trust from himself and thereby ceases or never begins to so believe, is doing something wrong. This is a substantive issue that I set aside.

phenomenon, the thought goes, can be theoretically significant if reference to the phenomenon is (at least sometimes) necessary to showing that something other than that phenomenon obtains, and if the theory is supposed to describe this other something. Such a phenomenon is necessary for the theory to do something else which the theorizer desires the theory to do, and so the phenomenon is theoretically significant. If trust theories, then, are supposed to explain how other-trust can be required and in fact is required, inconsistency will be theoretically significant if reference to both self-trust and the withholding of other-trust is necessary for showing that there are requirements to other-trust which are not being met.

With this sense of theoretical significance in hand, the second part of the objection seeks to show how reference to inconsistency may be necessary for showing subjects' failure to meet actual other-trust requirements. One might think that self-trust isn't merely an activity that subjects engage in when they do engage in it – subjects have *policies* to self-trust¹⁰. One might further think that if a subject has a policy to self-trust, then she implicitly takes some stance regarding reasons. According to the objection, then, this stance (which I will take to be an implicit *belief*) results in requirements to have a policy to similarly trust others, and so in some cases only by referring to the policy to self-trust can it be shown that a subject who fails to hold a similar other-trust policy is doing something wrong.

In response, I'm glad to concede the point of the first half of the objection. If inconsistency must be referenced to show that trust requirements apply and aren't met, then inconsistency has a place in a trust theory. The second half of the objection,

¹⁰ Thanks to Stan Husi and Edward Hinchman for, in various ways, providing the basis of this second part of the objection, and to the latter for pushing me on my more basic insignificance claim, leading to the first part of the objection.

however, fails to show that inconsistency is theoretically significant in such a way to outweigh theoretical insignificance. I'll consider two ways in which reasons might be implicit in a subject's having a policy to trust, and I'll argue that on either understanding of implicit reasons, inconsistency is theoretically insignificant, such that any significance is outweighed.

First, one might think that if a subject has the policy of trusting herself, then in virtue of holding this policy, she must also think there exists a reason why this policy is appropriate. This implicit stance or belief regards merely the *existence* of a reason, not the character or "content" of any such reason. The lack of a parallel policy to trust others would then result in inconsistency (or something very similar to inconsistency as I've analyzed it): the subject takes there to be a reason to have a policy to trust herself, so she's failing to treat as existing what she takes to exist (namely, a reason to have a policy to trust) in failing to form a policy to trust those similar to herself.

There are two main problems with this. First, the existence of some reason or other for a policy would seem to leave open the possibility (given what the subject believes, that is) of the unspecified reason being a practical reason, and as argued earlier, practical reasons aren't necessarily re-applicable in multiple cases. If a subject desires to avoid skepticism, for example, a policy to trust herself may be necessary for avoiding skepticism, but it's not clear that a policy to trust others in general is necessary to avoid skepticism. The way a subject treats herself can have an effect that her treating others in that same way doesn't have, and as a result, practical reasons for self-trust policies won't necessarily be re-applicable with regards to other-trust policies.

A second problem with this proposal is that, even if we set aside this first issue and limit the sort of reason that the subject implicitly takes to exist to epistemic reasons,

inconsistency will still turn out to be theoretically insignificant. For the sake of argument (at least), let's conceive of the putative reasons the subject implicitly takes to exist as being epistemic reasons: when a subject has a policy to self-trust, she implicitly takes there to be some epistemic reason to have a policy to trust herself. The question, then, is whether or not the implicit belief in (the existence of) an epistemic reason will be sufficient to generate inconsistency charges.

It would seem that for a subject to be failing to appropriately treat a putative reason as existing, she would have to think that the reason also exists when she is faced with the opportunity to have an other-trust policy. Or, since the existence belief is supposed to be implicit in the case of her self-trust policy, we may prefer to think that 1) at the very least the subject must believe, when faced other-trust opportunities, that certain conditions obtain with regards to another's faculties and environment, and 2) in addition to this belief in these conditions, it must also be the case that these conditions are somehow relevant. However, as the subject believes merely in the *existence* of a reason, there are *no* relevant conditions (with one exception, as I'll argue below). In other words, because the subject doesn't (implicitly) believe *about* any particular feature or set of features *X* of her faculties or environment that they provide a reason to have a policy to self-trust (by hypothesis), no feature or set of features of another subject's faculties or environment can be relevant.

At best, then, only if a subject takes another subject's faculties and environment to be identical (and not merely similar or very similar or nearly identical), will there be any relevance and therefore inconsistency.¹¹ If the second subject has identical faculties in an

11 I say "at best" because I'm not sure that any requirements arise without reflection on the subject's part when it's mere existence of a reason that a subject believes in.

identical environment, then, *whatever* the putative reason could be which is (implicitly) taken to exist for the first subject, that putative reason would also be present in the case of the second subject. There can be no relevant differences, because there are no differences, and because some reason is taken to exist and there are no relevant differences, the reason must hold in the second as well as the first case.

One might object that implicit reasons in holding a policy are necessarily generalizable. However, in the analysis above, reasons *are* generalizable: what a subject implicitly takes to be a reason, involving faculties X and environment Y (and resultant accuracy), is such that it applies to *any* subject with faculties X in environment Y. This is what would, if what I said above is correct, make identical faculties and identical environments sufficient for inconsistency charges in cases of failing to form an other-trust policy. However, if an objector wants there to be generalizability that extends beyond the subject's specific sorts of faculty and environment, it's not clear why such generalizability is necessary. It's not clear, in other words, why a subject's implicit putative reason *must* be generalizable beyond situations in which a subject has faculties X in environment Y. After all, the subject with the self-trust policy implicitly takes there to be a reason for her policy to trust *herself*. This reason, then, should relate her (sort of) situation to her accuracy (assuming we're dealing with an epistemic reason). But it's not obvious why such a reason should also relate (sorts of) situations different from her own to accuracy of beliefs.

If we take it that identical faculties and identical environments are almost never believed to exist by subjects, then inconsistency will in fact play very little role in a trust theory that is supposed to show actual other-trust requirements. This is another sort of theoretical insignificance, and so, while "in theory" reference to inconsistency may be

necessary to show wrongdoing, the sort of case in which inconsistency would need to be referenced would be too low to make the concept very useful.

This discussion may suggest a second way one might think reasons may be implicit in policy holding. Instead of a policy-holding subject implicitly taking some reason or other to exist for her to hold her policy, we might think that a policy-holding subject implicitly takes some *particular* (putative) reason X to exist. In other words, the subject implicitly believes *about* some features or set of features X of her faculties and environment provide a reason to have a policy of self-trust. This way of understanding a subject's implicit belief would block my theoretical insignificance response, as some particular feature or set of features of a subject's faculties and environment may be shared by other subjects' faculties and environment much more easily than *all* the features of a subject's faculties and environment.

This view about implicit beliefs seems plausible if we take our subject to have certain background beliefs not implicitly arising out of her self-trust policy, e.g. background beliefs about brains, favorable conditions, etc. In that case, though, it may be that our subject is guilty of some other failure if she fails to trust (either herself or others), given these background beliefs. However, if these background beliefs are insufficient on their own to generate some requirement, then this may be a way for policies to generate other-trust requirements. On the other hand, if our subject doesn't have any (relevant) background beliefs, then it seems somewhat mysterious how it might be that a subject implicitly takes X to be a reason, and equally mysterious why it is that the subject takes X versus Y or Z to be a reason.

Works Cited

Foley, Richard. 2001. *Intellectual Trust in Oneself and Others*. Cambridge University Press.

Foley, Richard. 2005. "Universal Intellectual Trust". *Episteme*. 2(1): 5-12.

Zagzebski, Linda. 2012. *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief*. Oxford University Press.