



Twenty-First Century Papers:
On-Line Working Papers from the
Center for 21st Century Studies
University of Wisconsin–Milwaukee

Judging as Inviting Self-Trust

Edward S. Hinchman

Twenty-First Century Papers: On-line Working Papers from the Center for 21st Century Studies * <http://www.21st.uwm.edu> *ISSN 1545-6161
*University of Wisconsin - Milwaukee * P.O. Box 413 * Milwaukee, WI 53201 * 414.229.4141 * 414.229.5964 (fax) *ctr21cs@uwm.edu

Director: Daniel J. Sherman

The Center for 21st Century Studies, a UW System Center of Excellence at the University of Wisconsin – Milwaukee campus, is a postdoctoral research institute founded in 1968 to foster crossdisciplinary research in the humanities. This series of occasional, on-line Working Papers provides a forum for rapid distribution of ideas in texts that are not yet ready or suitable for publication in more formal academic publications, but still offer valuable content. Usually the authors of Center Working Papers will be Center fellows, invited speakers, or others with significant ties to the Center, although we reserve the right to make exceptions. We regret that we are unable to consider unsolicited submissions.





Judging as Inviting Self-Trust*

[Draft of 20 November, 2007]

Edward S. Hinchman

hinchman@uwm.edu

Department of Philosophy
University of Wisconsin–Milwaukee

Twenty-First Century Papers: On-Line Working Papers
from
The Center for 21st Century Studies
University of Wisconsin–Milwaukee

Number 8

November 2007

© Copyright 2007 by the Center for 21st Century Studies
all rights reserved



* Thanks to Krista Lawlor and to participants in Susan Hahn's seminar at UWM, Spring 2006, for stimulating comments on an earlier draft.

What follows extends a project that I've been developing for several years on the role of trust, including self-trust, in rationality and personal autonomy. I'm interested equally in epistemic questions (what to believe) and in practical questions (how to act), though the present paper addresses only the former. In an earlier paper on the epistemic side of my project, "Telling as Inviting to Trust" (Philosophy and Phenomenological Research, May 2005), I offered an explanation of how you can get an entitlement to believe something simply by trusting someone when she tells you it's so, the entitlement deriving not merely from the assurance but from the trustworthy character of the one who gives it. The present paper adapts my explanation of believing-on-another's-word to belief as such, via the observation that you can't form a belief without trusting someone's judgment – if only your own. I view intrapersonal reasoning about what to believe as an internalization of the interpersonal cases I analyzed in that earlier paper. In each instance, belief-formation has a second-personal dynamic: an invitation to trust on one side, trusting acceptance of that invitation on the other. If I'm right, the dynamic at the heart of all belief manifests not autonomy but a trusting heteronomy, whereby you render yourself appropriately receptive to judgment – your own or another's. The idea that judging is inviting self-trust first occurred to me in 2002, but I didn't get a chance to develop it properly till 2005-06, my fellowship year at the Center for 21st Century Studies. I appreciate the support.

-o-o-o-o-o-o-o-

Judging is regarded by some as 'inner' assertion.¹ And asserting that p is regarded by some as representing yourself as knowing that p.² If we combine these theses, we make judging that p representing yourself 'inwardly' – that is, to yourself – as knowing that p. Though I do not accept either component thesis on its own,³ I believe that the combined thesis is close to correct. (As I'll explain in section VIII, to get the correct thesis we need merely replace 'knows' with 'has conclusive warrant.')

Though one may reject either or both of the component theses, I aim to reveal the philosophical attractions of putting them together.

These attractions follow from my emphasis on the intrapersonal relations at the core of judgment and belief. Judging manifests two distinct dimensions of self-reliance, I'll argue, and invites the relation of self-trust. Believing, I'll argue, is accepting the invitation. This explains

an interesting linguistic datum, that it makes sense to speak of ‘trusting your own judgment’ but not of ‘trusting your own belief.’⁴ And it provides the resources for a novel treatment of epistemic normativity.

I begin by developing a case that motivates my explanation of the linguistic datum (§§I-II), an explanation targeting the possibility of the species of doxastic akrasia that underlies it (§§III-V, XII). Having formulated my contribution in intuitive terms (§VI), I further argue that it offers an important corrective to both pragmatist and strong evidentialist accounts of doxastic norms (§VII). I then present my accounts of judgment and belief in detail (§§VIII-X), fending off some objections (§§XI, XIII) before sketching the more fundamental contribution to questions of epistemic normativity.

One way to conceptualize this contribution is in terms of the metaphor that belief has an ‘aim.’ Belief does not have an aim, on my account, but judgment does, and judgment’s aim can be revealingly specified as both belief and knowledge. Judgment aims at belief (and truth) by aiming at knowledge. Since judging is inviting self-trust, and the invitation is to trust *as a knower*, accepting the invitation is implicitly crediting yourself with knowledge. This knowledge account of judgment explains belief not merely as a state or as a status but as a robust trust relation.

I

Imagine you’re taking an autumnal solo hike on a trail that fallen leaves frequently obscure. The feeling creeps over you that you’re no longer on the trail. So you stop. You look around. Then,

after a minute or two, you form a judgment that the trail is “over there” to the right (as you put it to yourself, ostending with your gaze), and the judgment generates – or is simply manifested as – an inclination to resume your hike in that direction.

Now consider four variations on this scenario, in each of which we’ll naturally say that you either do not or should not trust your own judgment. First, imagine you’re a government surveyor just returned from several weeks in a national forest updating maps of its hiking trails. Since hikers rely on the maps you update, when ‘on the job’ you hold yourself to high epistemic standards and invest each judgment with far more care than you would if you were merely trying to keep from getting lost yourself. The problem, of course, is that you’re not on the job today, and epistemic stringency will yield no cartographic improvements. But you can’t shake yourself free from the habit of holding yourself to these exacting standards, even when you can see that the exacting standards are pointless. By your own lights you’d do better to make a call about where the trail continues and get on with your hike. But you can’t bring yourself to make that call.

Second, imagine you’re a student in your first philosophy course, taking a hike after class. You’ve just had your first encounter with Descartes’ evil genius and find yourself gripped by the possibility that you don’t know any of things you’d formerly assumed you knew. When you begin to feel you’ve lost the trail and find the doubt causing your legs to slow down, it feels like a piece of this more principled doubt and renders you unable to reach a decision how to proceed.

Third, imagine you’re giddy with delight at some news you’ve just received and find yourself disinclined to take any decision very seriously. “Lost? Well, then let me be lost!” you

exclaim as you find yourself inclined to charge off in the direction of your first hunch as to where the trail might be. (I'll develop this case in more realistic detail in section III.) Fourth, go to the other extreme and imagine you're morbidly depressed at the death of a loved one. Here you think, without the exclamation point, "Lost? Well, then let me be lost. I'm all loss anyway..." – as you find yourself inclined to shuffle off in the direction of a casually despairing stab at where the trail might be. Imagine that you do make a judgment, however pointillistic, in each of these last two cases, and consider them up to the point at which you're about to start moving.

The question raised by the first two variations is whether you should trust your faculty of judgment. You are unable to settle on a judgment, and the question is whether this doesn't show that your faculty of judgment is untrustworthy. The question raised by the second two variations is whether you should trust not only your faculty of judgment but an actual judgment that you've made. If we answer both questions in the negative, a third question arises: Does all this mistrust share a common basis?

We can easily imagine these cases such that you should trust yourself in none of them, and on a common basis: your status as an untrustworthy judge. In the first two cases, your judgmental faculty seems infected with a misplaced caution: you ought to be reaching a decision, yet you find yourself unable to. Moreover, you yourself can easily feel that the caution is misplaced: your judgmental faculty is untrustworthy from your own perspective. In the second two cases, you judge where perhaps you ought not to have judged, and again you're well positioned to see that the judgment should not be trusted. The first two cases needn't amount to mere hesitancy. We can imagine them raising the same questions of self-relation as the second

two cases. Moreover, this common basis explains why the untrustworthiness is manifest from your own point of view. Your judgment is not adequate to the demands of your own perspective. As the one who occupies that perspective, you therefore should not trust your judgment.

II

This diagnosis of the cases counters an alternative that begins from the thought that if you doubt your faculty of judgment, you simply can't make a judgment. That's what's going on in the first two cases, the alternative claims, and it shows that my diagnosis of the second two cases cannot be correct. The alternative explains the second conjunct of the linguistic datum as follows.⁵ You can't doubt your belief because (a) there is no special faculty of belief, since you use your faculty of judgment to form beliefs, and (b) when you believe that *p* you've already judged that *p*, so you *do* trust your faculty of judgment. On the alternative diagnosis, doubting a judgment that you've made simply reopens the deliberation by which you made it, thereby directly undermining the judgment. I'll say why I do not accept this diagnosis in sections III through V.

Let me first codify what I believe the cases show about the explanandum of any theory of judgment and belief. On my diagnosis of them, cases such as those just presented show that there can be a gap between judging and believing. The gap that appears here is not the gap revealed by the apparent possibility of believing contrary to your 'better' judgment. The gap revealed by my second two cases works the other way round: you judge without forming any belief on the matter. If in these cases you resist the inclination that would send you back on your hike, we can coherently say that though you judge that the trail is over there you do not believe it. My approach to belief and judgment aims to explain this gap.

Acknowledging the gap thus creates space for a distinctive species of doxastic akrasia. In this species, you judge that *p* but do not believe that *p*. In what I'll call the 'standard' species, by contrast, you both judge and believe that *p* – though both judgment and belief govern your conduct at a level beneath that of your reflective autonomy and stand in conflict with more considered judgments and beliefs.⁶ What distinguishes the two species is that in the first you do not trust your judgment and in the second you do. Each species of akrasia is puzzling in its distinctive way.

To see more fully how the species differ, imagine my second two cases as follows: you judge that the trail is over there, feel the inclination to resume your hike in that direction, and then reflect that your mood renders you untrustworthy on this matter but march off anyway, too elated or depressed to care. That would be crucially unlike the standard sort of doxastic akrasia, in which you manage to judge *and* believe that *p* by mechanisms that you do not and cannot acknowledge, an untrustworthiness that extends to your entire self-trust relation. In the newly imagined case three or four – the earlier version plus follow-through on the inclination – you don't trust yourself, but in the standard case of doxastic akrasia you do. In some standard cases it may be precisely this self-trust relation that you don't want to acknowledge. It thus seems possible to trust yourself in ways that you in some sense 'cannot' acknowledge – though it's admittedly hard to explain how the possibility works. The newly imagined case three or four is hard to explain in a different way. How could a judgment you don't trust control your behavior? That is, how could you both conclude a doxastic deliberation whether *p* by forming a disposition to act as if *p* (should the matter be relevant to action) without thereby 'buying into' this conclusion in the way characteristic of belief? This only superficially resembles the standard case of doxastic akrasia, in which through self-trust you do buy into your conclusions, albeit in a

way you cannot acknowledge. Each kind of case gives rise to puzzles, but the puzzles are very different.⁷ I'll sketch my account of the new species of akrasia in section XII.

III

Before proceeding to my account of judgment, which explains this non-standard species of akrasia, let me say more in defense of my claim that there can be the gap I've characterized between judgment and belief. One way to characterize judging and believing is to say, with Nishi Shah and David Velleman, that judging is the act of affirming that p and believing a consequent affirmative attitude toward p .⁸ If that's how you define the terms, then you'll have a hard time seeing how there could be an akratic gap between judging and believing. You'll be tempted, with Shah and Velleman, to regard the subject of my inquiry as 'ineffable.'⁹ For Shah and Velleman, drawing the doxastic conclusion that p is a matter of coming to regard p as true in a way that is governed, both normatively and descriptively, by the standard of truth. The key to the status of this act as a *judgment* – that is, as the forming of a *belief* – is this classificatory element in the subject's self-conception: she conceptualizes her cognitive act or attitude with a concept – judgment or belief – that is constitutively informed by the standard of truth. And this classificatory element works the same for judgment and for belief, the difference between them involving nothing more than the grammatical fact that one is an act and the other an attitude.

Shah and Velleman do acknowledge that, as they put it, "[a]rriving at the judgment that p doesn't necessarily settle the question whether one now believes it, since one may find oneself as yet unconvinced by one's own judgment."¹⁰ They consider a case in which the subject judges that the plane on which she is traveling won't crash but nonetheless finds herself believing that it

will. As they diagnose the case, an irrational phobia has “had a dominant hand” in determining what the subject believes.¹¹ By Shah and Velleman’s own account, however, it seems that this subject’s attitude toward the proposition that the plane will crash cannot be a belief, since her contrary judgment shows that the attitude is not governed in the right way by the standard of truth. True, she fears that the plane will crash, and that fear grips her so deeply that she cannot help acting on it as she might act on a hypothesis she just can’t shake. If she really does judge that that plane will not crash, then both her fear and her conduct – acting on and reasoning from the hypothesis (looking nervously at the exits, planning an underwater escape, and so forth) – are irrational. But we cannot say, with Shah and Velleman, that she irrationally believes contrary to her own judgment.

The species of doxastic akrasia that I’m examining does not amount to believing contrary to your own judgment. It amounts, more fundamentally, to *failing* to believe what you judge. So this species of akrasia is present in Shah and Velleman’s case, but not by virtue of the subject’s ‘belief’ that the plane will crash. It is present by virtue of the subject’s failure to believe that the plane won’t crash, given that she judges that it won’t crash. But we cannot explain the akrasia in the way that Shah and Velleman try to explain it – as her not being “convinced” by the judgment because a phobia holds her in its sway.

It is revealing that Shah and Velleman view the transition between judging that p and believing that p as a matter of ‘convincing’ yourself that p. It reveals that they view ‘conviction’ as, in effect, a synonym for ‘belief.’ The assimilation is natural enough – until we reflect that judging that p is *closing* deliberation on the question whether p, which does not leave open whether you are convinced that p. If you are not convinced that p, then your deliberation

whether p is *not* closed, so you haven't really judged that p . Your judgment does not leave open whether you are convinced that p ; it leaves open only whether you trust it – that is, whether you trust yourself qua judge – and manifest the deliberative closure in a settled commitment. It may be that fear prevents you from trusting yourself in this respect. Judging that p is not a matter of 'convincing' yourself that p , just as it is not a matter of 'telling' yourself that p . We can say that judging is *being* convinced – that is, undertaking a conviction. The only question left open is whether you trust that conviction.¹²

Of course, that formulation merely states my thesis in the language of 'conviction'; it is not an argument in its own right. We still need to know how it works. We need in particular to understand how judgment and belief stand in a relation that could break down in this way. Shah and Velleman's characterization of the transition from judgment to belief as "ineffable" suggests an assumption that there is nothing to understand here. I'll now demonstrate why the assumption is mistaken.

IV

For the purposes of my larger argument in this paper, I can agree with Shah and Velleman that concluding a doxastic deliberation by forming a judgment is a matter of regarding p as true in a way that is descriptively and normatively governed by the standard of truth. I can concede that if you've made a mistake about what the standard of correctness is for the kind of act you're performing, then you aren't really performing the kind of act you meant to be performing – judging – but instead doing something else. The problem is that Shah and Velleman overlook another kind of mistake you could make and thus could worry you're making: you could worry

you're applying the correct standard incorrectly. The worry might take either of two forms. You could worry that you're misweighing the evidence, or you could worry that you're failing to govern yourself by the correct standard for how much evidence is required for you to be entitled to conclude your deliberation with a judgment. You could of course have worries of both sorts at once, but I'll proceed by considering them in turn. (Even when you have both sorts of worry, it's easy enough to prize them apart.)

If you're worried you're misweighing the evidence – perhaps overlooking some evidence ready to hand, perhaps reasoning badly about evidence in hand – it seems your attitude towards your possible error has to be deliberative, and that this deliberative basis for the worry suffices to keep the deliberation open. Otherwise put, insofar as you're actively, occurrently worried that you've misweighed evidence, your deliberative stance has to be that the 'conclusion' you've reached is not a judgment but a hypothesis. And if you're merely disposed to such worry, then we can still say that you haven't really closed deliberation, since the disposition renders you disposed to keeping it open. This alternative therefore confirms Shah and Velleman's view. If the issue is merely that you've misweighed evidence, then there is no possible gap between judgment and belief. To 'doubt' your judgment on this basis could only be to doubt your faculty of judgment. And if you doubt your faculty of judgment, then you cannot form a judgment and must abandon any judgment you've made. Less grandiosely, such worries would typically manifest a more localized inability to close deliberation or to keep it closed. Either way, there is no judgment to mistrust in the akratic way introduced in section II.

Still, what if you're worried not that you've misweighed the evidence but that you're operating with a mistaken sense of *how much* evidence is required for you to count as entitled to

close deliberation with a judgment? We can codify the issue in terms of an ‘epistemic standard,’ construed as governing not primarily whether others may say of you that you know but whether you may treat available evidence as entitling you to close your doxastic deliberation. I’ll say much more about how the standard works, and about the relation between knowledge ascription and doxastic entitlement, in later sections. For now, I need only the idea that the entitlement to close doxastic deliberation with a judgment is sensitive to ‘what’s at stake’ in the context in which you’re deliberating.

If you’re worried that you’re operating with a mistaken sense of this standard, that’s importantly different from worrying that you’ve closed deliberation through misweighing evidence. If you’re worried that you’ve got the standard wrong, then you’re worried that you’ve closed deliberation when you should have left it open¹³ – not because you’ve made an error *within* the deliberation but because you’ve made an error *about* the deliberation. You’re worried that you’re operating with a mistaken sense of what’s at stake in your doxastic context. Again, we assume you’re governing yourself by the standard that constitutes what you’re doing as a doxastic judgment – you aren’t worried that you’ve mistaken *that* standard. You’re worried that you’re misapplying that standard in the doxastic context in which you find yourself. You’re clear that your judgment will be correct if and only if it is true, but the recognition doesn’t at all help you meet the standard. To meet the standard, the standard constitutive of judging that p as opposed to some other mode of accepting that p, you have to govern yourself by another standard determining how high a threshold your evidence must meet for you to be entitled to conclude that p. If you’re worried you’re operating with a mistaken sense of this latter standard, then *any* conclusion you reach will strike you as possibly illegitimate. You may be weighing the evidence in your possession perfectly properly, and you may not be overlooking any evidence

ready to hand. Still, perhaps you should have kept your mind open and devoted more time and energy to further inquiry.

Such a worry does not by itself reopen deliberation. There are three possible cases in this neighborhood, and yours is the one in the middle: (a) your worry is idle second-guessing, a thought that bounces around in your head, making you uneasy but without affecting your self-relation; (b) the second-guessing affects your self-relation by leading you to doubt the judgment you've reached, but without leading you to abandon that judgment and reopen deliberation; (c) self-satisfaction is restored when you align yourself with the worry, abandon the judgment and reopen deliberation. In the middle case, (b), you judge that *p* without (yet) believing that *p*.¹⁴

V

One might reply that there could be no such middle case, because the only way for your worry to move out of 'idle' into gear is for it to reopen deliberation. One might argue that even if the worry hasn't *yet* reopened deliberation, that is its telos: not distancing you from your own judgment but undermining the judgment. Perhaps it takes time for the worry to reach its goal, but once you embark on this path you are in the course of abandoning that judgment. Why should it make any difference that the process takes time?

It is true (if we wish to speak thus) that this is the 'telos' of the worry. Moreover, my claim that (b) is a distinct case isn't merely the claim that (b) codifies a temporal stage in the transition from (a) to (c). The problem with this argument is that it misconceives how the worry functions in driving you down that path. The transition from (a) to (c) isn't a mere intensifying

of the worry. In fact, the worry needn't grow more intense at all. The difference is that in case (a) you don't at all 'identify' with the worry: it merely floats around in your head. In case (c) you do identify with the worry, and as soon as you identify with the worry you no longer identify with the judgment that it undermines. A natural metaphor views you as 'switching sides' from the judgment to the worry. This process depends on a shift that is not driven by the worry, or by any other motive, merely on its own. After all, any motive can 'float around in your head,' in the sense that you feel it, without being or ever coming to be something with whose operations you identify your epistemic agency.

My theory of belief is in the end a theory of the 'identification' at the core of epistemic agency. But that piece won't fall into place till section XV. For now, I'll rest with the following reply to the argument that there is no stable middle case between the idle worry and one that undermines judgment. It is plausible, I concede, that the criterion of whether such a worry is merely idle is whether it plays any deliberative role. But it's a mistake to assume that the only deliberative role such a worry could play is to reopen the very deliberation about which you have these doubts. That's how the worry would operate if it were the first sort of worry I described: one directed at your weighing of evidence. But we're now imagining the second sort of worry – that you're operating with a mistaken sense of the epistemic standard – and that very different sort of worry gets a deliberative grip not by *reopening* deliberation but by opening a new deliberation: a deliberation on how to deliberate. This higher-order deliberation doesn't necessarily undermine the judgment you formed on the basis of your earlier first-order deliberation; it merely places you at a certain distance from it.

How should we characterize this distance? It's the distance, I propose, created by

mistrust in an epistemic authority. Mistrust in an epistemic authority needn't generate disagreement. When you mistrust a testifier, for example, you needn't view her as stating a falsehood. Your mistrust merely means that you can at most treat her assertion as *evidence* for the truth of what she says: 'She is right,' you can think, 'but I don't believe what she says merely because I take her word for it.' You can deliberate whether she has deliberated the matter rightly, but as long as you're wondering whether she has deliberated rightly you aren't simply taking her word for it – even if you believe that her conclusion is correct.¹⁵ Of course, when you mistrust your own judgment that p you can't treat it as mere evidence until you begin a new deliberation whether p, and if you're merely deliberating how to deliberate whether p you're leaving open that your earlier deliberation whether p was correct – in which case you won't redeliberate whether p. The point is that while you're deliberating how to deliberate whether p you don't trust your judgment that p – but without necessarily abandoning that judgment either. You neither trust your judgment that p nor abandon that judgment. In other words, you're in case (b): judging without (yet) believing.

What would such a worry look like in a realistic case? We could elaborate either of the two cases given schematically in section I. Take the first of these two (the third presented overall), in which your delighted mood leads you to treat your doxastic context as relatively undemanding. Feeling lost, you look around, backtrack a hundred feet to see if you recently missed a turn, then give the evidence available in perception and memory a few seconds' thought and draw a conclusion: 'The trail is over there.' You need have no worries about your competence in executing these tasks. But we can easily imagine you worried that your upbeat mood is leading you to misconceive your doxastic rights and obligations. 'Shouldn't I be taking this more seriously?' you non-idly worry. But you don't find yourself disposed to take it more

seriously. So there you are: both not taking it more seriously and mistrusting this fact about yourself.

If I'm right that such a worry, though non-idle, needn't force open your deliberation, we can see that there can be a gap between judgment and belief. Moreover, since with these remarks I've begun to talk about it, we can see that Shah and Velleman are wrong to characterize the manner in which you might succeed in bridging this gap as "ineffable." In the remainder of the paper I'll have plenty more to say about the self-trust relation whereby you pull off what Shah and Velleman rightly call this "perfectly familiar accomplishment."¹⁶ It's this very familiarity, I fear, that has bred a philosophical contempt for the transition's complexity. My guiding hypothesis, by contrast, is that our self-relations, like our interpersonal relations, tend to be fraught with philosophically fascinating dynamics just where we feel them to be most familiar.

VI

My account of judgment, which I'll introduce in this section and the next and elaborate more precisely in sections VIII and IX, provides the resources to explain how there could be an akratic gap between judgment and belief. (I'll complete the explanation in section XII.)

Judging, I've claimed, is (roughly¹⁷) representing yourself *to yourself* as knowing. In resuming your hike, then, you manifest this 'inner'-directed representation: that you take yourself to know that the trail is in that direction. It's the possible gap between the you doing the representing and the you to whom the representing is directed that explains the possibility of a

gap between judgment and belief. You can mistrust your judgment because you can mistrust your judging self.

If you didn't think you knew it, then you should have kept deliberating – treating your mental state as a mere hypothesis, not as a judgment. Moreover, self-consciously acknowledging that you don't know has the force of reopening the deliberation, as long as the issue remains a live one. If the question whether *p* poses a live issue for you, then when you take yourself not to know that *p*, you take the question whether *p* to stand in need of further deliberation, which is to treat your disposition to act as if *p* as a mere hypothesis. If you have reached a conclusion and thus take the matter not to require further deliberation, it follows that you take yourself to know the answer.

VII

This view has a number of attractions, I would argue more broadly, the most immediate of which is that it avoids the pitfalls of its competitors. Call 'pragmatism' the view that reflection on a purely practical consideration – whether it would promote some practical purpose if you believed the proposition – could serve to close deliberation. Call 'strong evidentialism' the view that only evidential considerations – indicating likelihood that the proposition is true – could play a role in closing deliberation. Neither view can be right, and my approach provides the necessary middle way between them.

Pragmatism must be false because, if the deliberation really is whether *p* and not merely whether it would be good to act as if *p*, no purely practical consideration addresses the question

you're deliberating. Setting aside cases in which the proposition has practical content, where a practical consideration can serve as evidence, the question whether p is simply not the same as the question whether it would promote some purpose if you believed that p . If you're asking whether p , it won't help your deliberation – at least, not directly¹⁸ – to answer this other question. But strong evidentialism must also be false, because if practical considerations didn't play *some* legitimate role in closing deliberation many if not most deliberations could never be legitimately closed. In many instances, for example, we close deliberation because we cannot afford to spend any more time deliberating. But the consideration that you cannot afford to spend any more time deliberating whether p is not relevant to the question whether p .¹⁹ And every deliberation is bounded by some such practical constraint – if only the exigencies imposed by your own mortality.²⁰ These practical considerations must play some role, though not an evidential role, in shaping the course of your deliberation and in setting conditions under which you may legitimately conclude it. The question is how they could do this.

This dialectic sets the framework for David Owens's *Reason without Freedom*.²¹ Owens doesn't draw a distinction between judgment and belief, but he argues that your belief that p must be motivated by a *non-reflective* impression that you have a conclusive warrant or ground for p , since (a) only the impression that you have a conclusive ground could close your deliberation and (b) a *reflective* appreciation of how pragmatic factors create a context in which the ground counts as conclusive could not figure in doxastic deliberation. Now I accept the spirit of Owens's (a), though as I'll explain I prefer to speak of 'the disposition to treat yourself as knowing' instead of 'the impression of a conclusive ground.' I do not, however, accept even the spirit of Owens's (b), his claim that the impression of a conclusive ground (or, as I would say, the disposition to treat yourself as knowing) must be non-reflective. I accommodate the

principal motivation for (b) by emphasizing how judgment is shaped by reliance on a non-*judgmental* appreciation of the contextually appropriate epistemic standard – non-judgmental, but not necessarily non-reflective. One central aim of this paper is to explain how there could be a non-judgmental species of self-reflection. I view this self-relation as posing a question of self-trust, resolving which manifests reflection not through positive deliberation or judgment but simply insofar as the resolution manifests a counterfactual sensitivity to evidence of your own *untrustworthiness* on the matter.

Trust, we can say generally, is a species of willed dependence, where the dependence is willed insofar as it is under appropriate guidance of a counterfactual sensitivity to evidence of untrustworthiness in the trusted. ‘Appropriate guidance’ here means: you would not trust if you had evidence that the trusted is not worthy of your trust. Trust need therefore involve no active assessing at all, if the trusted is trustworthy; it is enough merely to be (and, counterfactually, have been) disposed to notice and respond appropriately to evidence of untrustworthiness should there be (or have been) any. Though you may of course have practical reasons for or against trusting, I don’t think the *epistemic* reasonability of trust derives from your possession of epistemic reasons or warrant *for* the trust. Self-trust thus manifests a relation of self-reflection, though it is non-deliberative and non-judgmental.

Returning to Owens’s (a), then, we can see that we lack a rationale for positing a non-reflective ‘impression’ of warrant at the basis of judgment. Looking ahead, we’ll see that my account improves on Owens’s both by elaborating the self-relations at the heart of these epistemic attitudes and by locating the basis of judgment in a self-reflective, albeit non-deliberative, attitude toward a state of your own character.

From the perspective that I share with Owens, the question again is how practical considerations can play a role, though not an evidential role, in shaping the course of your deliberation and in setting conditions under which you may legitimately conclude it. My account of judgment explains how they could do this by noting that a practical consideration could play a role in the deliberation whether p by determining, not whether p , but whether you may legitimately conclude that p . It could do this if (i) you may legitimately conclude that p if and only if you know that p , and (ii) a practical consideration can play a role in determining whether you know that p (or at least in determining whether you have conclusive warrant for p).²² Those who reject strict invariantism grant that practical considerations – or, at least, considerations that are not evidential in the above-defined sense²³ – can play a role in determining how stringent an epistemic standard governs a knowledge-ascription in the relevant context. One of my aims in this paper is to treat that concession as the key to an account of judgment, whereby judging is representing yourself to yourself as having met the contextually appropriate epistemic standard. If practical considerations can play a role in determining that standard, but without thereby registering in how you deliberatively represent yourself (to yourself) as meeting the standard, then they can play a legitimate role in closing deliberation.

Though my account rejects both pragmatism and strong evidentialism, it aims to capture what pragmatists and evidentialists on their opposing sides rightly emphasize. Pragmatists are right that judgment and belief are to an important degree governed by practical interests. And evidentialists are right that these interests cannot play a direct role in producing judgment or belief. Both sides are right, because judgment and belief are expressions of your self-ascription of knowledge, and because practical considerations play an important role in determining whether you in context count as knowing. Only evidential considerations – indicating likelihood

of truth – count in answering the question whether *p*, but doxastic deliberation is governed by a norm that says to conclude that *p* when but only when you know that *p*, and non-evidential considerations – indicating what is practically at stake in the relevant context – do count in answering the question whether you know that *p*. You aren't deliberating *whether* you know that *p*. (You could, of course, deliberate that question, though you needn't be and typically aren't deliberating it.) Nonetheless, the norm governing deliberation is identical to the norm governing the ascription of that concept. You judge that *p* by aiming to satisfy the application conditions of 'I know that *p*.'

Again note well: I am not claiming that judging that *p* is judging that you know that *p*. The condition under which you may conclude that *p* is the condition of your knowing that *p* (or at least, as I'll clarify in the next section, of your having conclusive warrant for *p*). As I'll argue, this is the condition of your knowing that *p* (or of your having conclusive warrant for *p*) relative to the epistemic standard operative in the context of a possible *assessor* of the claim that you know that *p*. If you – optionally – *also* judge that you know that *p* and assess that claim from your own perspective, then you occupy all three positions yourself: subject, ascriber and assessor. But I'm going to argue that this case is not the most fundamental. My most general thesis is that the self-relations at the heart of judgment and belief internalize relations that we bear most fundamentally to others.

VIII

Let me now lay out my accounts of judgment and belief in detail.

Since it's simpler, I'll first state my account of belief. Call being disposed to act as if p 'accepting' that p. By 'act' I include mental acts as well as physical acts: the disposition in question may in some cases be restricted to thinking or feeling as if p.²⁴ Now:

(B) S believes that p iff S accepts: that she judges that p and in so judging knows that p.

Note that, on (B), believing that p directly includes being disposed to act as if you judge that p. But it indirectly includes judging that p, since you wouldn't be disposed to treat yourself as knowing that p unless you judged that p.²⁵ I don't mean that believing that p presupposes a discrete act of judging that p; I mean merely that believing that p includes the cognitive dispositions that figure in judging that p. It adds to those dispositions the further disposition to act as if you know that p.

To get my account of judgment, we need to amend my earlier formulations and eschew the factive 'knows that p' in favor of the non-factive 'has conclusive epistemic warrant for p.' By 'conclusive epistemic warrant for p' I mean: warrant sufficient for knowing that p, on the assumption that p. 'S knows that p' is factive, since to assert that S knows that p is (*inter alia*) to assert that p. To assert that S has conclusive warrant for p is not yet to assert that p. As I intend them, there is no other difference between the assertions: 'has conclusive warrant' does nothing more than subtract factivity from 'knows.' In light of the gap between judging and believing, we cannot construe judging in factive terms because judging that p is compatible with not trusting yourself whether p. We need an account of judging that p that does not commit the subject to accepting that p. My proposal is therefore not, strictly, a 'knowledge account' of judgment but a 'conclusive warrant account.'²⁶

Here is the full account I'll defend:

(J) S judges that p iff S accepts:

- (i) that she has conclusive epistemic warrant for p ,
- (ii) that she accepts (i), and
- (iii) that (i) because (ii).

According to (J), judging that p is accepting that you've achieved an epistemic status – having conclusive warrant for p – as a result of your accepting that you accept that you've achieved this status.²⁷ Accepting that you have conclusive warrant for p , and that you accept that you accept this, is compatible with not accepting that p . You manifest the first acceptances by, for example, being disposed to reply 'P' (or 'I'd have to say that p ') to the question 'P or not p ?' and, most fundamentally, by no longer being disposed to deliberate the question whether p – that is, by being disposed to treat the matter as closed (no longer considering evidence for not- p , etc.). You can coherently have these dispositions while lacking the more general disposition to act as if p . As a self-doubting hiker you may treat the question of the trail's location as closed without being disposed to resume your hike (though you are disposed to resume your hike in the direction of the trail). You are not disposed to resume your hike in any direction, while being disposed to treat yourself as having concluded where to resume it, because you are not disposed to treat your concluding of this matter as authoritative. Though you of course *may* (and typically do) accept that p when you judge that p , you needn't do so. Sometimes you judge that p without accepting that p , sometimes you judge that p while accepting that p but without believing that p , and sometimes you do all three.

IX

Each account can be stated in a sentence – as (J) and (B) – but much remains to be said about how they work. Let me devote two sections to elaborating my explanatory framework, before turning to some natural objections in section XI.

The first explanatory challenge is simple: How could you have conclusive warrant for a proposition *because* you accept that you do? A way to express (J) is to say that judging is inviting warrant-generating self-trust. Though I have parried some objections in previous sections (and will consider further objections in sections XI and XIII), I must explain more fully how self-trust could generate warrant. In this section I'll characterize the self-relation in more detail; in the next I'll say how it can generate warrant.

Consider what (J) says. You judge that the trail is 'over there,' it says, not merely by accepting that – i.e. being disposed to act as if – you have conclusive warrant for the proposition. You're disposed to act as if you have such warrant when pretending that the trail is over there, but when you pretend that something's the case you don't (except rarely) judge that it is. A distinctive species of *commitment* to *p* figures in judging that *p* that does not figure in imagining, supposing or pretending that *p*. These other acts embody other species of commitment, but not the one distinctive of judging. What, then, is the latter species of commitment? It doesn't help to say that it's commitment to the truth of *p*, since supposition and pretense embody their own species of commitment to the truth of *p*. When you suppose or pretend that *p*, you suppose or pretend that *p* is true, and in that respect you're – imaginatively, we might say – committed to the truth of *p*. It doesn't help, therefore, to note that judging aims at the truth. So, in its distinctive way, does every other cognitive propositional attitude and act.²⁸ The question is *how* judging

aims at the truth.

To this question (J) provides the answer. In judging you aim at the truth, it says, by aiming to accept that p only when your acceptance that you have conclusive warrant for p gives you conclusive warrant for p . Consider in turn the two clauses of this necessary condition, (ii) and (iii) in (J). To say, with clause (ii), that you accept that you accept that you have conclusive warrant for p is to say that you're disposed to act not only as if that's true (i.e. that you have the warrant) but as if you're disposed to act as if it's true. This requires merely a minimal degree of self-consciousness – not that you think about yourself or your dispositions but that you're disposed to act as if you have those dispositions. (And remember: being disposed to act as if you have conclusive warrant for p – that is, to treat yourself as having the warrant – needn't involve the disposition to act as if p . Perhaps you don't trust the former disposition.) Clause (ii) merely describes the self-consciousness involved in the reflexivity of human thinking, which needn't yet involve self-awareness.²⁹

The reflexivity of judging does not yet distinguish judging or believing that p from imagining, supposing or pretending that p . None of these acts is a mere disposition to act as if p . Sometimes we *are* merely disposed to act as if p . If stricken with an obsessive-compulsive disorder you may be disposed to act as if your hands are always dirty, not merely washing them repeatedly but thinking about their uncleanliness and as a result feeling a vague unease. But that's not yet to judge or believe that your hands are dirty. It's not even to imagine, suppose, or pretend that they're dirty, in the agential senses under discussion here. Of course, it's possible that as a result of these dispositions you do come to imagine, to suppose, or even (by seeking out evidence) to believe that your hands are dirty. But the ailment itself is merely to be stricken with

uneasy thoughts of uncleanliness. Though it may cause you to imagine or even to believe that your hands are dirty, the ailment itself, a mere disposition, does not include the latter reflexive dispositions. By contrast, each of the cognitive acts and attitudes – imagining, supposing, pretending, as well as judging and believing – entails that you are disposed to act as if *p*, insofar as you are, in part because you're (in the right way) disposed to act as if you're disposed to act as if *p*.

Clause (ii) involves more, of course, than mere reflexivity. Call a disposition to act as if you're disposed to act as if *p* a *reflexive disposition* to act as if *p*. Clause (ii) asserts that you have a reflexive disposition not to act as if *p* but to act as if you have conclusive warrant for *p*. But that more complex reflexive disposition also fails to distinguish judging and believing from other cognitive attitudes. When you imagine that *p*, you're imaginatively disposed to act as if you have conclusive warrant for *p*. (For example, you imagine yourself entitled to tell others that *p*, which is to imagine yourself possessed of the relevant authority.) When you suppose that *p*, you're suppositionally disposed to act as if you have conclusive warrant for *p*. And so on.³⁰ Again, each cognitive attitude embodies its distinctive species of commitment to *p*. The commitment is merely the distinctive way in which you're reflexively disposed to act as if you have conclusive warrant.

The differentia of judging lies instead in clause (iii), which says that to count as judging that *p* you must be disposed to act as if your disposition to act as if you have conclusive warrant for *p* gives you conclusive warrant for *p*. No other cognitive act or attitude embodies that implication. When you merely imagine or suppose that *p*, you don't presume that your dispositions actually *give* you conclusive warrant for *p*; you merely manifest a disposition to act

as if you possess it. But how could an act or attitude satisfy this differentia? Here we get to the heart of the matter. How could your reflexive disposition to act as if p give you conclusive warrant for p ?

If it's this self-reflexive disposition to act as if you have warrant for p that gives you warrant for p , then it's a fact about you that gives you warrant for p . And if in judging you represent yourself as warranted on the basis of this fact about you, then in judging you represent yourself as warranted on the basis of self-reliance. This is not the insane thesis that you bootstrapped your way into possession of the warrant. It's merely the tautology that when judging you can't avoid relying on your own status as reliable in judging.

Since parallel tautologies are equally true of the other acts, we must say specifically what reliability in judging comes to. We can again contrast judging with imagining, and it helps to begin with an interpersonal case of each. Say you're telling tall tales to an audience that you hope to impress with your inventive wit. You get your imaginative powers in gear and then see where they take you. Here you have to rely on your dispositions as much as in the judgmental case, though the propositions that you express in telling your tales are not propositions that you judge. The difference is that while your imaginative enterprise does require the exercise of certain virtues – virtues of inventiveness – it does not in itself require exercise of the virtue exercised in good judgment. Yes, you have to be a good judge relative to the propositions that you judge – e.g. that your audience will react in this way or that – but you don't have to be a good judge with respect to the propositions that you imagine. Your self-reliance in each case is a confidence that the reflexive disposition to act as if p manifests the specifically proper virtue.

In parallel fashion, I propose, your reflexive disposition to act as if you have conclusive

warrant for p could give you conclusive warrant for p if the disposition manifested the virtue proper to judging. It is on your possession of this virtue that you distinctively rely when you judge.

X

What, then, *is* the virtue proper to judging? The temptation is to leave it at: reliability in judging the truth. Of course, such reliability has got to figure in the virtue, since to the extent that you're not reliable in judging the truth on a subject matter you ought not (other things equal) to be willing to judge. But I don't think that's all or even the main part of it. The most fundamental part of the virtue is your ability to hold yourself to contextually appropriate epistemic standards, where – as I've suggested and will now argue – these are the standards that someone whom you might tell what you judge would appropriately use in assessing whether your speech act is an expression of knowledge.

When you judge you present yourself as accepting the truth and as having met the epistemic standard appropriate to the context in which you judge. There are thus two facets to the virtue proper to judging: reliability in judging the truth and reliability in sensing the contextually appropriate epistemic standard. We can say that these are two separate virtues, if we want, and that your self-reliance when you judge involves a confidence in your exercise of both. But there's a reason to resist this separation maneuver, and it derives from my basis for claiming that reliability in sensing standards is more fundamental to judging than reliability in judging the truth.

The basis for the claim is that it is at best misleading to speak of ‘reliability’ in judging the truth simpliciter, since there is no conceivable human purpose served by relying on someone as a judge of truth simpliciter. Say you know of someone only that he makes mostly true judgments. Why should this matter to you? After all, the true judgments may be restricted to matters of no conceivable human concern. You yourself could be very ‘reliable’ indeed if you restricted your judgments to extremely trivial matters. It is at best misleading to speak of ‘reliability’ where there is no real chance of anyone actually relying on you. What matters is not ‘reliability’ in judging the truth simpliciter but reliability in judging the truth on relevant subject matters taken in context. But what subject matters? what contexts? My hypothesis is that you can’t reliably answer these questions in a given case if you aren’t reliable in sensing the contextually appropriate epistemic standard.

To illustrate, let’s return to the hiking case and reflect on how unclear it is in advance what would *count* as looking for a trail in such a context. And reflect on how broader issues affect the definition of your practical problem. Sometimes you find a trail by following your eyes, sometimes your ears, sometimes your nose. Reliability in which of these faculties would make you reliable here? And in which specific ways of exercising them? But if it’s not clear in advance what counts as looking, it cannot be clear in advance what counts as reliability in finding. We thus discover a new version of *the generality problem*.³¹

We can solve the generality problem in theory by insisting that the virtue on which you rely when you judge include a skill in solving the generality problem in practice. The appearance of circularity is both inevitable and non-paradoxical. If you could not in practice solve the generality problem, by relying on your ‘sense’ of what it takes to have conclusive

warrant for p in context, you could never come to judge that p . Note that the skill is exercised in second intention. To judge, you needn't actually have the skill. (That *would* be paradoxical!)

You need merely be reflexively disposed to act as if you have it.

Why not obviate the challenge and attempt to get by with nothing but hypotheses? You can't get through life forming nothing but hypotheses, I suggest, because the interpersonal contexts in which you find yourself require that you sometimes commit yourself to others, giving them 'your word.' What is this 'word' that you give? It is tempting to view it as the product of your judgment, and then to view telling – that is: informing, letting know – as the externalization of this inner attitude of judgment. I suggest we view matters the other way round. Your 'word' is indeed the product of your judgment, but judging is the internalization of such telling: it is premised on the possibility of letting a possible interlocutor know – if not an actual then an imagined or at least a posited interlocutor. That you can judge that p when aware that you can't claim to know that p to your actual interlocutors merely shows that you needn't conceive of relevant interlocutory space as defined by the actual.

Back in the woods, you know you can't be absolutely certain you're not making a mistake – both about where the trail is and about what standard determines conclusive warrant in judging where the trail is. But when you judge that the trail is 'over there' and thereby resume your hike in that direction, you treat yourself as possessing a complex interpretive skill. Not only do you take yourself to be conclusively warranted in judging where the trail is, but you take yourself to know *how* to be thus warranted by recognizing the appropriate epistemic standard. This recognitional capacity is the virtue proper to judging.

Your grasp of the relevant standard is potentially as complex as your grasp of how the

present context fits into a broader stretch of the life you share with possible interlocutors (including your own future self) who may (perhaps in memory) rely on your judgment (though such self-reliance is not in other respects like testimony, and I doubt you could coherently conceive of a life that you shared in this respect *only* with your future selves). Of course, the broader grasp does not by itself provide a determinate context. Your sense of this broader context both constrains your sense of the contexts you encounter moment by moment and is their product. To be capable of judgment is to think of yourself as engaged in both of these enterprises, with each informed by the other. Though it must include an appropriate species of purely truth-conducive reliability, the virtue proper to judgment – that on which you rely when you trust your judgment – is more fundamentally hermeneutic. The virtuous judge is reliable both as a source of truth and as an interpreter of what that requires.

XI

Before returning to doxastic akrasia in section XII, let me reply to some objections. (I'll consider yet another objection in section XIII.)

First, it's natural to demur that when you judge that the trail is 'over there' you don't take yourself to *know* that the trail is in that direction: what you take yourself to know is at most that the trail is *probably* in that direction, where the 'probably' marks your subjective confidence in the proposition. This objection raises large issues,³² but perhaps I can step round them for now simply by clarifying what case we're imagining. In the case we're imagining, you weren't trying to figure out the probability that the trail is in this direction or that. You were trying to figure out *where* the trail is – or, as you'll put it, *whether* the trail is 'over there.' So if you manage to close

this deliberation, the conclusion you draw concerns *where the trail is*, not merely the probability that it is there. (We'd get a different case, but with exactly the same complexities, if we imagine you deliberating the probability that the trail is in this or that direction.) Now we're imagining that you do close your deliberation, by which I mean that you draw a conclusion, not merely that you cease to deliberate (which of course is compatible with giving up on the question). You don't, we're imagining, judge merely that the trail *may* – with such-and-such probability – be in that direction, i.e. that in light of this probability it's worth walking over to look for further evidence of the trail. If that were your attitude, then your walk would manifest the status of your deliberation as ongoing rather than closed. That could of course have been your attitude; you could have decided not to resume your hike but to continue your gathering of evidence in that direction. In the case we're imagining, however, you judge that the trail is in that direction – 'over there' – and then on that basis resume your hike. I propose that when you close your deliberation, you take yourself to know that the trail is 'over there,' in this direction in which you are now inclined to set out.

Second, a worry about (J)'s reference to 'conclusive warrant.' I claim that when you judge that p you are disposed to act as if your being disposed to act as if you have conclusive warrant for p gives you conclusive warrant for p . That formulation makes it seem that, by your lights, the basis of your warrant lies in your own dispositions – which again is an insane thesis (unless, of course, the judgment is about your dispositions). My idea, however, is not that you accept your disposition to act as if you have conclusive warrant for p as the basis of your warrant for p but that you accept it as an epistemic criterion of that warrant, by which I mean *as authoritative* for whether you have it. How could your disposition to act as if you have conclusive warrant for p count as such an authority? There are two dimensions to this authority:

first, the disposition must be informed by the doxastic-deliberative weighing of evidence and other warrant-providing considerations; second, the disposition must have been formed through your having concluded deliberation only when these considerations provided you with warrant that is *conclusive*. I codify the second dimension of authority as clause (iii) of (J): what you treat as ‘giving’ you the conclusive warrant is your disposition to treat the considerations that you’ve weighed as sufficing for conclusive warrant on the matter – that is, crucially, in the *present context* of deliberation. Of course, again, the basis for this warrant lies in the considerations themselves. But you treat the warrant *as conclusive* by concluding your deliberation. As I’ll go on to argue at length, this reliance on your own capacity appropriately to conclude your deliberation is the self-relation at the heart of judgment. It is, I’ll argue, the same as reliance on your capacity to gauge what’s at stake for an implicit audience of potential interlocutors.³³

Third, a worry about my gloss of ‘accepting that p’ as being in a complex dispositional state.³⁴ This entails that judging that p is most fundamentally a state rather than an act; where it is an act, it is the act of entering into the dispositional state. The gloss may seem at odds with my characterization of judging that p as concluding the deliberation whether p, and generally with my use of agential notions in discussing these issues. In reply let me clarify that I’m offering a quasi-reductive account of judgment: *reductive* because I reduce the agential aspect of judging to non-agential dispositions, but only *quasi* because I do not aim to reduce the concept of judging. My account does not reduce the concept of judging because my analysis (J) appeals to warrant, and we cannot explicate the concept of warrant without explicating the concept of knowledge, which we cannot do without relying on the concept of judging.

Though my analysis of judging, (J), reduces judging to complex dispositional states, my

argument for (J) focuses on the subspecies of *agential* judging, the judging that happens when you perform the act of concluding a deliberation. This is merely for ease of exposition. I hold that you deliberate whenever you let evidence and other warrant-generating considerations weigh with you and that you conclude a deliberation whenever you satisfy the conditions of (J). It's easier to think about what's going on when you satisfy the conditions of (J) as the termination of a self-conscious process. I use that instance to clarify why we should think both (a) that the process can unfold unselfconsciously and therefore (at least, given certain popular accounts of agency) non-agentially, and (b) that (contra David Owens) we should not infer from this fact that terminating deliberation fails to manifest a species of self-reflection. Though concluding a deliberation is not (given certain common assumptions about agency) an act that you perform, neither is it merely something that happens to you. In adjudicating the self-relation at the heart of (J), you manifest a capacity for self-criticism that brings judgment into the realm of what you take responsibility for, even if the way in which you take responsibility for it is not by performing any action.

Let me emphasize that if you really don't think you *know* that the trail is over there, you can refrain from judging that it is without rendering yourself motionless and doomed. You can form the hypothesis that the trail is over there, and then continue your deliberation by walking over there to confirm whether you were right. If that's your attitude toward the proposition, then you shouldn't tell your companion that the trail is over there but merely that it *may* be, or that you suspect that it is. (In either case you can specify your degree of confidence.) You have a choice between shifting modality or substituting a different proposition, most naturally one not directly about the trail but about this deliberation that you've not yet closed. Whichever way you go, the case is different from the one we're discussing. We are not discussing a case in which

you don't judge but merely hypothesize, or in which your judgment is merely *about* your deliberation.

XII

How then do I account for the 'non-standard' species of doxastic akrasia introduced in section II? The key is that (B) and (J) allow this possible case: you are disposed to act as if you have conclusive epistemic warrant for p , and to treat that disposition as what gives you the warrant, without believing that p . For you to count as believing that p more is required than your thus satisfying the conditions of judgment; you must also be disposed to act as if these complex cognitive dispositions constitute you as a knower. Belief can come apart from judgment because it is possible to doubt that your disposition to act as if you have conclusive warrant for p gives you conclusive warrant for p – even when you are disposed to treat it as giving you the warrant. This is possible because it is possible for your dispositions concerning your own epistemic entitlements to come apart from your dispositions to act. When you doubt your judgment that p , you typically won't be disposed to act as if p . Acting as if p is not impossible in that condition, however – though if you do act as if p while doubting your judgment that p , it will be a self-dissociated or self-alienated species of action. Insofar as what you do presupposes epistemic entitlement to p , you won't be able to recognize yourself in your doings. Though *your judgment* is in control of your conduct, you yourself are not.³⁵

This self-alienation manifests your doubt in the reliability of the faculty whose exercise informs your sense of epistemic entitlement: your faculty of acting as if you have conclusive warrant for a proposition only when you have conclusive warrant for that proposition. To

understand the possibility that I have in mind, we need to be clear that this is not the ('standard') sort of doxastic akrasia in which you both believe that p and believe that not- p (or otherwise believe in a 'motivated' way that you cannot acknowledge). In that case, you might doubt the reliability of your faculty for determining when you have conclusive warrant simply because it has enmired you in this conflict. But such self-doubt plays no role in the akrasia itself: the problem in standard akrasia isn't that you don't trust your judgment. In the non-standard case under discussion, by contrast, that is the problem: you judge that p (by being disposed to treat the fact that you're disposed to act as if you have conclusive warrant for p as giving you conclusive warrant for p) while nonetheless at the same time doubting your judgment (by not being disposed to act as if p). How could you do this? By doubting the reliability of the dispositions that perform the job of judgment, that of closing the deliberation whether p only when you have conclusive warrant for p . If you go on to act on these dispositions, you aren't in on your action in the way you are when you believe. This species of doxastic akrasia would lead to action, or at any rate behavior, that is governed by judgment but not by belief.

This stance may resemble the self-relation prescribed by Pyrrhonian skeptics. But such self-doubt cannot be simply a broad sense of your unreliability on the question whether p , since if you have such a broad doubt in your reliability you'll not have managed to form the disposition at the heart of judging that p . The self-doubt must therefore be more specific.

Closing deliberation only when you have conclusive warrant for p involves gauging not merely whether p but what's at stake in the context of deliberation in a very particular respect. As I've argued, you must gauge whether a possible audience to whom you could go on to express this judgment would thereby come to have a testimony-based warrant for believing that p

themselves. The mental act of closing deliberation in a judgment expresses, in effect, a commitment to the cogency of possible speech acts that you could perform on its basis. If, having judged that p , you were to tell someone A in circumstances C that are relevantly like your own, would that telling (other things equal) allow A to know that p (or at least to possess conclusive warrant for p)? Well, what epistemic standard holds in C ? The standard that would allow A , this possible interlocutor, to know in C . But that would have to be the standard that A would aim to meet in judging that p herself in C . By the same reasoning, that standard would have to be the standard prevailing in the contexts of A 's possible interlocutors. In each case, in sum, the standard tracks the epistemic needs of those who would assess the act of telling in the course of wondering whether to rely on the teller's word.

You have to go by your 'sense,' as I've been putting it, of this contextual standard, and this 'sense' obviously can't amount to another judgment, on pain of a vicious regress. Your 'sensing' of this standard is simply your accepting that (i.e. your being disposed to act as if) your having closed deliberation by judging that p gives you conclusive warrant for p – relative to the *second-personally* relevant epistemic standard. If you judge that p while doubting your ability to sense this standard, your self-dissociation goes deeper than traditional forms of skepticism, whether Pyrrhonian or Cartesian, can make intelligible. Such self-doubt would manifest a dissociated sense of where you stand with possible interlocutors – both real and imagined. The self-alienation thereby expresses the role of possible interlocutors in constituting your sense of self. You'd better hope the self-doubt doesn't generalize.

The hypothetical speech act is an implicit claim to know (or at least to possess conclusive warrant) relative to the standards of pertinent interlocutors needing the information you thereby

impart.³⁶ And your mental act of judging is committing yourself to the cogency of performing such a speech act. You may judge though fail to believe because you may doubt your ability in this instance to gauge the cogency of inviting others to judge on your say-so.

XIII

Can't you intelligibly tell someone 'I believe that p, but I don't know that p'? Indeed you can. Properly understood, such cases are not counterexamples to my accounts of judgment and belief but confirm them.

Here is an example from Seymour Hersh: "Richard Armitage, the Deputy Secretary of State in Bush's first term, told me, 'I think Iran has a secret nuclear-weapons program – I believe it, but I don't know it.'"³⁷ In saying that he doesn't know that Iran has a secret nuclear-weapons program, Armitage is acknowledging that the epistemic standard for assessing such a claim, especially as made by the recent Deputy Secretary of State, is very high. The disavowal of knowledge makes sense only because of this feature of the background. Compare: "Richard Armitage, looking out his office window at the torrents of rain pummeling the glass, told me, 'I think it's raining out – I believe it, but I don't know it.'" This could make sense only if Armitage and his interlocutor are engaging in some epistemology, with a skeptical challenge looming in the background. But then the disavowal of knowledge would play the same role as before in acknowledging the unusually high epistemic standard. Since the standard is unusually high, the self-ascription of belief in each case has a hypothetical air: I am disposed to act as if p, and I would count myself as believing that p if the epistemic standard weren't so unusually high. Armitage's 'I believe it' amounts only to a hypothesis plus an indication that the standard is

unusually high.

My account can thus accommodate cases in which it makes sense to self-ascribe belief – that is, to *say* ‘I believe’ – without self-ascribing knowledge. What you’d be saying is that you *would* know, and therefore feel entitled to judge, if the epistemic standard in your interlocutory context weren’t so unusually high. It’s an acknowledgment not only that but how the epistemic standard governing judgment is sensitive to context: since you’re telling *me*, and my doxastic context imposes a higher epistemic standard, you don’t feel entitled simply to tell me what you would feel entitled to judge relative to the needs of other possible interlocutors.³⁸

You needn’t, therefore, put yourself in position to let all your interlocutors know that *p* when you judge that *p*. But you can’t judge that *p* while regarding yourself as unable to let *any possible* interlocutor know that *p*. If you couldn’t let any possible interlocutor know that *p*, then you couldn’t (even potentially) give anyone your word that *p*. If you couldn’t (even potentially) give anyone your word, then you haven’t a ‘word’ to give. If you haven’t a ‘word’ to give – if you don’t regard yourself as epistemically capable of informing anyone on the question whether *p* – then you haven’t formed a judgment whether *p*. By putting the point in terms of your ‘epistemic’ capacities I mean to abstract from your physical capacities or your psychological capacities. You can of course judge that *p* while acknowledging that you can’t tell anyone that *p* because you’re gagged or because you’ve been given a drug that induces great shyness. Again, call the possible people to whom you physically and psychologically can or could speak your ‘possible interlocutors.’ My point is that if you are unable to give any possible interlocutor your word that *p*, then you haven’t really judged that *p*. As you’ll put it when they ask you whether *p*: ‘I just can’t say,’ meaning not that you haven’t any words to utter but that you haven’t a word –

your word, your assurance – to give.

XIV

Let me put my account in perspective before elaborating, in the final section, its core position on epistemic normativity.

My approach can be viewed as developing an idea of Edward Craig's: that the point of the concept of knowledge is to indicate potential informants, where an informant is someone who can let her audience know that *p* simply by telling them that *p*.³⁹ If the concept of knowledge serves to identify potential informants, the standard that a knowledge claim must meet is the standard governing the context of someone wondering whether the claim is true – since it is they who stand ready to be informed.⁴⁰ And if the norm you must meet to count as knowing is sensitive to the context of someone wondering whether you know, then in aiming to know you must aim to meet the standard governing that context (or those contexts, if there is more than one such person). Insofar as we can make sense of the thesis, this is the 'aim' that governs the formation of belief. (I'll say how the metaphor is potentially misleading in the next section.) My contribution here is twofold: first, we can make some sense of the idea that there is an aim governing and giving content to the attitude of belief; second, the aim is not merely truth but knowledge – thus second-personally construed. We aim at truth *by* aiming at knowledge.⁴¹

The aim governs belief by informing the more basic attitude of judgment that, as my title puts it, 'invites' belief. Strictly speaking, then, it is judging, not believing, that has an aim. Judging is concluding doxastic deliberation when you sense you have met the relevant epistemic

standard, something you can do only within a fundamentally second-personal intrapersonal relation. My approach is distinctive in treating not only knowing but also judging and believing as second-personal concepts.⁴² Just as telling is inviting your interlocutor to trust you,⁴³ judging is inviting self-trust. Believing is then taking up the invitation: accepting your own judgment as another might accept your testimony. Your status as knowing is thus what would enable *another* to know by trusting you exactly as you trust yourself in believing. And that other would know in just the same way. The first-person self-relation thus aims both to mimic and to warrant the second-person relation. Believing aims at knowing by instituting an intrapersonal facsimile of the interpersonal relation whereby you can enable another to know.

XV

My approach to judgment and belief makes room for a novel approach to questions of epistemic normativity. When you judge that p you form a reflexive disposition both to act as if you have conclusive warrant for p and to treat the disposition as giving you that warrant. As I've explained, I don't think you thereby count as believing that p – at least not insofar as believing figures in self-government.⁴⁴ For judging is pure disposition, albeit a complex one, and self-government is that region of our thought wherein we endorse or otherwise 'buy into' our dispositions – not merely acquiesce in them. (B) says what it is to buy into your judgment: it's to be disposed to treat your judgment as constituting you as a knower. By my account, the self-government constitutive of belief-formation is a disposition to treat yourself as factively entitled to a proposition by certain of your dispositions. The entitling dispositions are those in (J). In one respect, then, belief is a disposition. In another respect, however, belief is not a disposition.

Since it is a disposition not only to present yourself as knowing (as in judgment) but to treat the presentation as accurate, belief embodies a *claim* of knowledge that judgment only aspires to.

It is in that claim that you constitute yourself as subject to epistemic norms, since without such a claim you would not be a unified target for epistemic praise and blame. A mere disposition is something with which you must cope, but it is not yet anything for which you can be praised or blamed. Only a claim – as we say: a stance, a position – presents a target for praise or blame. Here you deserve praise or blame for being disposed to treat your dispositions as constituting you as a knower.

The self-reliance that you need in order to judge thus differs from the self-trust that you need in order to treat what you judge as a guide. If trust is not mere reliance,⁴⁵ then only the self-relation at the core of believing is a self-trust relation. Judging thus involves intrapersonal relations in one sense but not in another. The relations are to yourself – to your dispositions – but they are not to yourself qua person. It's only when you go on to believe, trusting your judgment, that you realize a bilateral intrapersonal relation, a relation in which you count as your own epistemic authority. Judging does invite that relation, projecting an audience whose epistemic needs it presumes to anticipate. But the final word rests with the one who would treat that presumption as amounting to knowledge. When you trust your judgment, you make an implicit claim that the presumption is warranted. When you act on judgment without self-trust, by contrast, what you do is in this respect self-dissociated: you do not credit your own presumption.

My accounts of judgment and belief in (J) and (B) are quasi-reductive, but the account of knowledge on which they ultimately rest is not. I view knowledge as setting the norm for

judgment and belief, and I suspect that the most we can say about this norm in general is how it functions. I hold that it functions primarily as what ensures that the results of your doxastic deliberations are useful not only to yourself but also to others. I follow Edward Craig in viewing the concept of knowledge as constitutively tied to the concept of an informant, but I don't think Craig has adequately explicated what the status of informant involves. On my view, being an informant involves more than being a reliable guide to the truth, since a good informant must be trustworthy in another dimension as well. When you let your audience know that p, you thereby assure them not only of the truth of p but of what it would take for them to know that p. You invite them to rely on you in believing that p in the way that you thereby show you rely on yourself.⁴⁶

The presumption informing this invitation manifests the aim of judgment, which in turn lends its content to the 'aim' of belief – in that now-widespread *façon de parler*.⁴⁷ On my account, belief does not literally have an aim, but it is not wrong to say that judging aims at knowledge. More exactly, when you judge that p you aim to institute a self-relation wherein self-trust on the question whether p expresses a more fundamental status of knowing that p. Though you can't count as knowing till you've instituted the self-relation, you presume that the judgment puts you in position to know simply through trust.

In an important respect, then, the 'aiming' metaphor is misleading. When you aim at a target and shoot, the target plays no agential role in determining whether your aim was correct. But when you aim to constitute yourself as a knower, the self thus constituted plays a role in its constitution. The role is target-like in its passivity but untarget-like in its agential bearing. When you trust yourself and believe, what you add to judgment's invitation is a reasonable

responsiveness to the trustworthiness of your own judgment. You needn't *judge* that you're trustworthy – which would in turn raise all the same questions about that new content. The responsiveness is not itself a judgment, a belief, or any other assessment. The agency you thereby manifest is not active but passive. It is reasonable, and thus agential, insofar as this passivity is informed by a counterfactual sensitivity to evidence of untrustworthiness: as I explained in section VII, you must be disposed *not* to trust insofar as there is evidence that your judging self is not worthy of it (and have been so disposed had that self not been worthy).

Judging is, in sum, less like aiming at a target than like leading on a dance-floor.⁴⁸

Success is not merely a status but a relationship.

¹ "The soul when it thinks is simply conversing with itself, asking itself questions and answering, affirming and denying... So I define one's thinking as one's speaking and one's thoughts as speech that one has had not with someone else or aloud but in silence with oneself" (Plato, *Theaetetus*, 189e-190a). "Judgment is to grasping a thought as assertion is to the expression of a thought" (Michael Dummett, *Frege: Philosophy of Language* [Cambridge: Harvard University Press, second edition, 1981], 298). See also Wilfrid Sellars, "Language as Thought and as Communication," *Philosophy and Phenomenological Research* 29 (1969).

² See Peter Unger, *Ignorance: A Case for Scepticism* (Oxford: Oxford University Press, 1975), Chapter VI; Michael Slote, "Assertion and Belief," in J. Dancy (ed.), *Papers on Language and Logic* (Keele: Keele University Library, 1979); Timothy Williamson, *Knowledge and Its Limits* (Oxford: Oxford University Press, 2000), Chapter 11; Keith DeRose, "Assertion, Knowledge, and Context," *Philosophical Review* 111 (2002). As he notes in footnote 6 (252), Williamson argues for a somewhat more general thesis from which my formulation follows.

³ One virtue of my elaboration of the combined thesis is that it generates a principled basis for rejecting the thesis as applied to assertion (i.e. the second thesis on its own). What some philosophers incorrectly view as a feature of assertion is in fact a feature of 'inner' assertion or judgment. I pursue this argument in [Author's paper 6].

⁴ Citing Wittgenstein's observation that "One can mistrust one's own senses, but not one's own belief" (*Philosophical Investigations*, trans. G.E.M. Anscombe (Oxford: Blackwell, 1956), p. 190), Richard Moran argues that the unintelligibility of 'mistrusting one's own belief' shows that "neither trust nor mistrust has any application here" (*Authority and Estrangement* [Princeton: Princeton University Press, 2001], 75). My aim in this paper is both to explain why that claim is true, in the respect that it is, and to argue that it is not true in the respect in which Moran appears to intend it. While Moran acknowledges that "insofar as I lack confidence in my judgment about X, I have no settled belief about it" (77), he does not appear to recognize the extent to which its constitutive relation to judgment implicates belief in robust self-trust relations just as such. One implication of my discussion is that Moran should have cast his treatment of doxastic authority and self-estrangement as an inquiry into the dynamics of judging as inviting self-trust.

⁵ I should note that it is not my aim to capture every aspect of ordinary usage of the terms 'judge' and 'believe.' Both 'I trust my judgment' and 'I don't trust my judgment' (along with second- and third-person variants) are common in colloquial English, but the general form 'I judge that p' (in any person) comes off as rather stilted, the verb 'to judge' having been relegated to specialized usage (as what an official Judge does, e.g. in law or

sport). Nonetheless, ‘I (don’t) trust my belief’ (or a variant) is not merely stilted but a solecism.

⁶ This formulation assumes that doxastic akrasia involves conflicted believing, an assumption rejected by Alfred R. Mele, *Self-Deception Unmasked* (Princeton: Princeton University Press, 2001), especially Chapter 4. But I could reformulate to avoid this assumption. My point is merely that in ‘standard’ doxastic akrasia the subject’s judgment does not come apart from her belief.

⁷ For example, though you may be ‘alienated’ from your behavior (or from the motives that produce it) in both cases, there is an extra dimension of alienation in the non-standard case: the dimension that I codify as your lack of self-trust. It’s a good question whether such alienated behavior should count as ‘full blown’ action (in David Velleman’s sense; see his Introduction to *The Possibility of Practical Reason* [Oxford: Oxford University Press, 2000]), but that’s the subject of a different paper.

⁸ Nishi Shah and J. David Velleman, “Doxastic Deliberation,” *Philosophical Review* 114 (October 2005), 497-534, at 503.

⁹ “A problem might be thought to arise in the transition from [the] act of affirmation to an affirmative attitude, but to our minds, there is no problem about that transition... Exactly how one accomplishes the transition is of course ineffable, but it is a perfectly familiar accomplishment, in which a proposition is occurrently presented as true in such a way as to stick in the mind, lastingly so represented” (*ibid.*).

¹⁰ “Doxastic Deliberation,” 507.

¹¹ “Doxastic Deliberation,” 508.

¹² If you have the intuition that conviction must be belief – that is, must manifest *trust* in your judgment – I simply ask you to set that intuition aside for now. I am putting my thesis in these terms only to commensurate it with Shah and Velleman’s discussion. Nothing turns on whether a conviction is a judgment or a belief. In my idiolect, it makes sense to wonder whether to trust my own conviction (but not whether to trust my own belief). But I concede that it also makes sense to say that when I doubt my judgment I am not yet convinced on the matter. I concede that ‘being convinced’ that p is somewhat ambiguous as between judging that p and believing that p.

¹³ You could of course also worry that you’ve left deliberation open when you should have closed it – in other words, that the standard is not as demanding as you’re taking it to be. But in that case you won’t have formed a judgment, so there won’t be anything (apart from the faculty of judgment) to mistrust. So we couldn’t get a gap between judgment and belief. This doxastic gap can appear only when you worry that the standard is more demanding than you’re taking it to be.

¹⁴ I add the ‘yet’ as a reminder that you don’t *dis*believe that p, if that means a deliberate refusal to believe that p. You merely lack the belief that p.

¹⁵ In these remarks I draw on the position in the epistemology of testimony that I defend at length in [Author’s paper 1] and [Author’s paper 2].

¹⁶ “Doxastic Deliberation,” 503.

¹⁷ Again, in section VIII I’ll replace ‘knows’ with ‘has conclusive warrant.’

¹⁸ It may, of course, help indirectly by helping you to diagnose prejudices skewing your deliberation whether p. Still, that’s obviously a different deliberation, and you could never conclude the first deliberation by appeal to such a consideration.

¹⁹ Unless, again, the proposition has practical content – for example, that you will soon stop deliberating. Throughout this discussion I’m going to set such cases aside, since these practical considerations have evidential content in this context and are not therefore *purely* practical.

²⁰ Even apriori reflections are potentially open-ended insofar as it isn’t an apriori matter how much checking and rechecking of your reasoning may be required to safeguard against mistakes. Sometimes you don’t know that p because your apriori reasoning was too slipshod to yield knowledge as opposed to a good guess.

²¹ (London: Routledge, 2000). I accept the broad thrust of Owens’s polemic in the early chapters of the book, though the very different account that I offer of *how* belief (via judgment) aims at knowledge departs from his on the role of reflection in closing doxastic deliberation. (Owens hasn’t the resources to consider my account, since he doesn’t distinguish judgment from belief.) In light of this difference, unlike Owens I am not in the business of denying autonomy or freedom its traditional role at the core of doxastic reason. I’m instead trying to reconceptualize what autonomy or freedom must *be* in order to play that role.

²² Once more, I’ll refine the formulations in section VIII, replacing ‘knows’ with ‘has conclusive

warrant.’ The lack of refinement shouldn’t matter in the meantime.

²³ I’ll assume that all such considerations are in some sense ‘practical.’

²⁴ My usage of ‘accept’ resembles David Velleman’s in “On the Aim of Belief” (in *The Possibility of Practical Reason*, 250, except that I’ve eschewed Velleman’s use of the proto-judgmental term ‘regard’ in my gloss of acceptance. Velleman holds that to accept that *p* is to regard *p* as true, where this ‘regard’ is the attitude held in common among believing, imagining, supposing, and pretending that *p*. My appeal to a disposition to act as if *p* captures this commonality, given that I include mental acts. (Velleman comes very close to my dispositional formulation when he characterizes the motivational role of belief on 255.)

²⁵ When you merely pretend to judge that *p*, you are disposed to act as if you judge that *p* without actually judging that *p*. But it is incoherent to be disposed both to act as if you judge that *p* and to treat this judgment as constituting you as a knower without actually judging that *p*. When you *merely* pretend to judge that *p* you necessarily don’t view yourself as thereby knowing that *p*. (You can, of course, believe that *p* without having the belief that you judge (or have judged) that *p* – just as you can believe that *p* without believing that you believe that *p*.)

²⁶ Another advantage of switching to a conclusive warrant account is that we thereby allow that a doxastic agent may blamelessly judge a false – though conclusively warranted – proposition. ‘But you didn’t have conclusive warrant for *p*’ is always a legitimate (pro tanto) criticism of someone who judges that *p*, whereas ‘But *p* is not true’ may not be a legitimate criticism where the agent has conclusive warrant for *p*.

²⁷ Note the formal parallel with the Gricean theory of meaning, on which producing speaker’s meaning is intending that your hearer produce a certain result as a result of his recognizing your intention that he produce this result. Indeed, I’ve modeled my account on Grice’s, transforming the interpersonal into the intrapersonal. (The issues of speaker’s meaning and judgment are obviously very different, however. All I’ve taken from Grice is the form of the account.)

²⁸ Velleman also makes the point that imagination, supposition and pretense ‘aim’ at the truth in the sense that they involve imagining-, supposing-, and pretending-true (“On the Aim of Belief,” *op. cit.*, 247-250).

²⁹ I do not in the least deny the possibility of self-deceptive or otherwise less than fully self-aware judgment or belief. It may indeed surprise you to find yourself judging that *p*, or to learn that you believe that *p*. You can be surprised by your own judgments because you can be deceived about what you judge. The reflexivity of judging has relevance not to whether you’re surprised but to what you’re surprised at – that is, to the *target* of your surprise: you’re surprised that you are disposed to act not only as if *p* but as if you are disposed to act as if *p* (where the proposition in question may be *you have conclusive warrant for p* or merely *p*). You are surprised not merely to find yourself disposed but to find yourself reflexively implicated in your dispositions.

³⁰ It may seem more accurate to say that you’re disposed to imaginatively (or suppositionally, etc.) act as if *p*, rather than that you’re imaginatively (suppositionally, etc.) disposed to act as if *p*. (Again, the proposition in question may be *you have conclusive warrant for p* or merely *p*.) And that may seem to provide a basis for distinguishing judging and believing from these other acts: when you believe, you are disposed to *act* as if *p*, not merely to ‘believingly’ act as if *p*. The problem, however, is that ‘imaginatively’ acting as if *p* is acting as if *p*. Or if we say it isn’t, we’ll have to explain how imagining differs from believing, which leaves us back where we started.

³¹ I mean the ‘problem of generality’ for reliabilism discussed in Alvin Goldman, ‘What is Justified Belief?’ in George Pappas (ed.), *Justification and Knowledge: New Studies in Epistemology* (Dordrecht: D. Reidel, 1979), Section II, and developed in Richard Feldman, ‘Reliability and Justification,’ *Monist* 68 (1985). More recently, see Earl Conee and Richard Feldman, ‘The Generality Problem for Reliabilism,’ *Philosophical Studies* 89 (1998).

³² The demurral may embody the mistake about the point of making a judgment that I mentioned in the previous section: it may rest on the idea that you could get through life forming nothing but hypotheses. I pursue this issue more fully in [Author’s paper 5].

³³ We’re in the vicinity of the worries about ‘bootstrapping’ that I assuage by a similar move in the account of intention that I defend in [Author’s paper 7] and [Author’s paper 8]. My point in both instances is that what you rely on is not merely the considerations that warrant you or give you reasons but how you weigh them.

³⁴ I’ll clarify other aspects of my dispositional analysis in section X, including the reflexivity built into (J).

³⁵ This formulation looks ahead to my account in section XV of what this possibility shows about epistemic normativity. In brief, what it shows is that, though judgment *invites* an intra-personal relation, the judge's self-relation becomes intra-*personal* only when the invitation is accepted in belief.

³⁶ I'll say more about the identity of these interlocutors in section XIII. It is crucial that they needn't include just anyone who happens to hear or otherwise witness your speech act.

³⁷ Seymour Hersh, "The Iran Plans," *The New Yorker*, 17 April 2006.

³⁸ As I explain in [Author's paper 2], when S's context is less demanding than A's it is possible for S to judge that p without being entitled to tell A that p. Moreover, when S's context is more demanding than A's, she may sincerely tell A that p without judging that p herself. Such cases reveal that it is not a necessary condition on sincere assertion that the speaker says what she believes. They also directly falsify the 'knowledge account' of assertion, since such an S may be perfectly entitled to assert that p, though in failing to believe that p she fails to know that p. (For the knowledge account of assertion, see the works cited in note 2 above. I pursue this point in [Author's paper 5].)

³⁹ *Knowledge and the State of Nature* (Oxford: Oxford University Press, 1990), especially Chapter V. I claim that this is the point of the concept of knowledge, not of testimony. Again, I don't claim that testimony necessarily passes on knowledge.

⁴⁰ John MacFarlane has developed this position most fully; see his "The Assessment Sensitivity of Knowledge Attributions," in Tamar Szabo Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology* (Oxford: Oxford University Press, 2005). I dissent, however, from MacFarlane's core defense of his position (see my [Author's paper 5]).

⁴¹ I agree here with David Owens (*Reason without Freedom*, 35). I do not, however, place any *explanatory* weight on the idea that belief has an 'aim,' so my account avoids the pitfalls that Owens lays out for the thesis that belief has an aim ("Does Belief have an Aim?" *Philosophical Studies* 115 (2003), 283-305).

⁴² 'Second-person concepts' may seem to echo Stephen Darwall's emphasis on 'second-personal reasons' in the practical context (see his *The Second-Person Standpoint* (Cambridge: Harvard University Press, 2006), an emphasis that he argues cannot be sustained in the epistemic context (55-60). As I clarify at length in [Author's paper 1], which develops the analysis of testimonial telling that I presented in [Author's paper 2] (and defends it against some recently published objections), my conception of the second-personal is quite different from Darwall's. For Darwall, second-personal considerations concern what is *owed* to someone (e.g. respect), whereas for me they concern what someone may *need* (e.g., as we'll see, a trustworthy informant). In a purely practical context, my dissent from Darwall's conception of the second-personal most fundamentally addresses his claim that reasons of (sometimes much needed) *care* are third-, not second-, personal (126-130). I pursue the practical issue in [Author's paper 3] and [Author's paper 4].

⁴³ See my [Author's paper 2].

⁴⁴ Some philosophers view belief in a way that doesn't let issues of self-government – broadly, the 'ethics' of belief – arise. It should be clear that I'm setting such approaches aside. Such approaches are typically not interested in distinguishing belief (or judgment) from other cognitive attitudes such as supposition and pretense.

⁴⁵ Annette Baier (in the papers on trust collected in her *Moral Prejudices* (Cambridge: Harvard University Press, 1994)) and Karen Jones ("Trust as an Affective Attitude," *Ethics* 107 (1996)) both rightly emphasize the difference between trust and mere reliance. I present my own reasons for emphasizing the distinction in [Author's paper 3].

⁴⁶ The show can of course be merely that: a pretense. Let me emphasize once again that I am not analyzing telling or asserting in terms of letting know, nor am I appealing to any comparison with letting know in my explanation of judging. 'S is letting A know that p' is factive in two dimensions: it entails that p, and it entails that S believes that p. Judging needn't involve a parallel factivity in either dimension: you can judge what is false, and you can fail to believe what you judge. The parallel between the intrapersonal and the interpersonal that I am trying to explain is only imperfectly (but still, I think, usefully) captured by the parallel between judging and letting know.

⁴⁷ Besides the work of Velleman and Shah cited earlier, see Peter Railton, "Truth, Reason and the Regulation of Belief," *Philosophical Issues* 5 (1994), 71-93, and "On the Hypothetical and Non-Hypothetical in Reasoning about Belief and Action," in G. Cullity and B. Gaut (eds.), *Ethics and Practical Reason* (Oxford: Oxford University Press, 1997). For a very different approach to the relation between aiming at truth and aiming at

knowledge from the one I've defended, see Bernard Williams, *Descartes: The Project of Pure Inquiry* (London: Penguin, 1978), 38-45, and Ralph Wedgwood, "The Aim of Belief," *Philosophical Perspectives* 16 (2002), section 6. Other approaches include Christopher Peacocke, *Being Known* (Oxford: Oxford University Press, 1999), 34, and Timothy Williamson, *Knowledge and Its Limits* (Oxford: Oxford University Press, 2000), 208.

⁴⁸ Expert dancers inform me that following an expert lead on the dance-floor requires a certain suppression of one's counterfactual sensitivity to evidence of untrustworthiness. But some such sensitivity is nonetheless in play: certain moves would rightly not be tolerated. And there is room to explain the appropriate degree of sensitivity in terms of what's at stake in the relevant context. (For an elaboration of the dance-floor case in illustration of the authority of *practical* judgment, see [Author's paper 8].)