

Knowledge, Noise, and Curve-fitting:
A methodological argument for JTB?

[to appear in R. Borges, C. de Almeida, & P. Klein (Eds.), *Explaining knowledge: new essays on the Gettier problem*. Oxford University Press.]

Jonathan M. Weinberg
University of Arizona

Theorizing what knowledge is given that it *isn't* simply justified true belief has been, of course, a major project of the last half-century, and even when authors haven't taken on any sort of conceptual analysis of knowledge, the need to motivate, and then accommodate, counterexamples to JTB remains nearly universal in epistemology. David Lewis includes Gettier on a short list of two possible cases in which conclusive philosophical refutations may have occurred -- the other case being Gödel (1983, x). The existence of this volume is itself further evidence of the disciplinary centrality of Gettier's work in epistemology, if any such further evidence were somehow needed! But there has long been a small but interesting set of epistemological dissenters as well (see, e.g., Kaplan (1985); Sartwell (1991); Hetherington (2012); and also Turri (2012) for a friendly if ultimately critical take on Gettier-dissent). I am here aligned with the forces of dissent, and will contend that we ought not take the rejection of JTB as a firm constraint on our theories of knowledge.

I. Noise, models, and overfitting: A general argument for simpler theories in philosophy

My argument for re-opening the books on JTB is fundamentally methodological, and starts with the method of cases -- short vignettes and our evaluations of their epistemic content, and a strong norm in practice that theories of knowledge must cleave fairly closely to those evaluations. Although general theoretical concerns can and do inform our discussions the inadequacy of JTB as an account of knowledge, nonetheless, I suspect those principles themselves are often backed, perhaps implicitly, by the case evidence. Some have contended that Gettier does not merely monger cases, but offers arguments for his suggested evaluations of them (e.g., Deutsch (2015); Cappelen (2012)). Whether this is true as a matter of proper historical understanding of the original paper -- though I have my doubts, and see, e.g., Weatherson's (2014) response to Cappelen -- it is clearly false about the way that Gettier's insights have percolated through epistemology. Quite clearly the focus in almost all discussions are about "Gettier cases" or "Gettier's counterexample" and so on. (For example, try googling on "Gettier case" or "Gettier cases", and then "Gettier's argument", and the difference in the number of instances is stark.) And much of the literature has been focused on generating more such cases, testing novel accounts against them, and repeating that cycle, to the point that quite a taxonomy of different Gettier cases can now be reviewed. (see Blouw et al., this volume). I take it that such case-based methods depend on the idea that we are on the whole better at tracking what is or is not a case of

knowledge, than we are at directly recognizing correct or incorrect principles about knowledge as such. We humans are at least middlingly decent epistemometers, when it comes to cases and especially cases within the ordinary reach of our experiences, and that's why it makes sense for us to appeal to our sense of the right verdict in such cases when we are building and critiquing our theories of knowledge. This is not at all to be committed to the extravagant view that we *only* can, should, or do appeal to such cases; obviously none of that is so. But to the fairly large extent that we do appeal to them in epistemology, we are making that presupposition of the broad reliability of our particularist capacity for epistemic evaluation. So our methodological discussion can take as a starting point that our epistemometric capacities are not by and large a matter of explicit inference or argument, and that it will generally not make sense to ask in regard to these verdicts, what an individual's deliberately-deployed premises and mode of inference was in arriving at it. In that way, they play a methodological role very similar to basic perceptual reports, whether there or are not any deeper similarities to be found there.

A quick terminological point: I will speak throughout of "verdicts" and "verdict data" and the like, so as to set aside the rather thorny debates about what intuitions are or aren't or even if we should speak in terms of "intuitions" at all. Not that such debates aren't important, but I do think they are orthogonal to the kinds of issues I will be attending to here. Also, I will be understanding justification here in the traditional sense of a status that can be had even by false beliefs, and thus JTB is something that is in conflict with the Gettier verdict data as standardly understood, in taking those cases to be ones in which justification is fully present but not knowledge. My main interest is in trying to explore ways in which inferential concerns can trump the apparent deliverances of our data. So, I will not be engaging here with versions of JTB that say the agents in Gettier cases do not know, but because they are not even justified in the first place, since such theories generate no such data-versus-inference conflict.

Similarly, although one could, at least in principle, attempt a methodological restoration of the prospects for JTB by attacking the basic reliability of the verdict data itself, I will not be pursuing such an argument here. I bring up this argumentative strategy only in order to set it aside. I have no wish to challenge that presupposition of decent-enough baseline reliability about cases, which has been defended amply in recent years (e.g., Williamson (2008), Boyd and Nagel (2014)). One immediate concern about such an approach in this context, for example, is that it will be too costly: if we are to throw out the verdict data wholesale, then how do we make a positive case for JTB, or any other theory of knowledge? Some may be surprised to see me granting such a claim of substantial baseline reliability to our verdict data, since the contrary has sometimes been attributed to me (e.g., by Nagel in her (2012)). But in fact I have long been happy to endorse it (e.g., Weinberg (2007); Swain et al. (2008); Alexander and Weinberg (2014)).

Granting this baseline reliability does not, however, not go very far towards

settling the pressing methodological questions about how best to go about deploying such verdicts in our philosophizing. There has been some attention to what further demands we might place on our practices with sources of evidence, beyond their being merely baseline reliable, and whether our practices in philosophy perhaps fall short of those demands (Weinberg (2007); Alexander & Weinberg (2014)). These discussions have all been focused on the conditions under which we make use of the verdict evidence. But comparatively very little attention has been paid, however, to the structure, nature, and demands of our *modes of inference* that take that verdict evidence as a source of premises. And that is where I will draw the materials from for my argument. Part of my goal here is to shine some light on the methodology of philosophical inferences, beyond the standing debates about philosophical sources of evidence.

For the susceptibility of our verdicts to error has methodological consequences beyond problematizing our initial deployments of them. This *noisiness* of this body of evidence also has implications for what inferences we might look to draw from them, and in particular, for when considerations of *simplicity* should require us to disregard some of the apparent deliverances of the verdict data.¹ There are already a handful of very interesting extant attempts to bring such simplicity into philosophical methodology, and I will now discuss them briefly, with a bit of an aim towards criticism, but more to set up a contrast with my own preferred approach. (They are all more or less consistent with each other, I think, but each brings very different sorts of background assumptions with it.) There are at least 4 different ways that some sort of appeal to simplicity can be introduced, in order to motivate overriding Gettier intuitions in favor of JTB: a *metaphysical* appeal to features of natural properties; a *metaphilosophical* appeal to the desiderata on conceptual explications a la Carnap; an *epistemological* appeal to holistic virtues of good theories; and, finally, what I will be deploying here, a *methodological* appeal to the principles of model selection.

Brian Weatherson in his (2003) works within a Lewisian/Siderian framework of reference magnetism, where the reference of a term is fixed largely but not exactly by use. The set of cases that a speaker or community of speakers applies the term sets a major constraint upon, but does not by itself precisely fix, the reference of the term. Rather, the referent is whatever property best maximizes both closeness of fit to those cases, and its naturalness as a property. Weatherson floats the idea that maybe JTB -- as the conjunction of what, he claims with at least *prima facie* plausibility, are three highly natural properties -- is so natural a property, that even if our intuitions have a Gettier-shaped bend in them, nonetheless "knows" in our mouths may really pick out JTB and not any sort of

¹ I am focusing on simplicity here, but another way in which noisiness can make trouble for philosophical inferences is that our modes of inference are often high in *epistemic demandingness* (Nado 2015) -- the theoretical claims in play tolerate very little variance from their highly specific predictions. I also pursue this concern in my (forthcoming), but not with as specific an eye towards Gettier and JTB.

JTB-plus-further-machinery-to-rule-out-Gettier-cases. (I will henceforth use "JTBG" as a shorthand for theories with such a structure.) More generally, the simplicity of an analysis should be expected, at least roughly, to track naturalness of property, and thus we have one *pro tanto* reason to prefer simpler theories.

Stephen Crowley and I raise a number of worries for Weatherson's account in our (2009), such as the possibility that even more natural than JTB is, simply, knowledge. In general, natural language terms of significant philosophical interest, like "knows" or, say, "cause", may be so eligible as referents, that we will be unable to use considerations of naturalness to motivate any sort of decompositional analysis whatsoever. (A kind of Zagzebski/Williamson-meets-Fodor kind of result.) But I want to focus instead on a different worry that Crowley and I put forward: it may be too methodologically intractable a criterion of simplicity to be of any real use in our philosophical practices. Given a few different rival hypotheses as the referent of "knowledge", it might just not be humanly possible to discern which is more natural than which. Consider, e.g., JUSTIFIED TRUE BELIEF and TRUE BELIEF PRODUCED BY A RELIABLE BELIEF-PRODUCING MECHANISM. It is, to put it mildly, not obvious which of those is more natural (and maybe they are equally so). It is not only not obvious now, it is moreover hard to see what further investigations we could do in order to better ascertain their comparative naturalness. And the problem only gets much, much worse if we suppose that we are faced not only with a larger set of hypotheses, but also a choice where various naturalness/closeness-of-fit trade-offs are under consideration.

A similar worry, I fear, would apply to any attempt to appeal to holistic virtues of JTB understood as a Carnapian explication of the concept "knows", as has been pressed recently by Olson (2015). On Olson's proposal, the pragmatic nature of explicatory projects provide the motivation for the holistic considerations, such as simplicity, fruitfulness, and exactness, that may ultimately give us reasons to override rather than solve the Gettier problem. One drawback of appealing to this Carnapian metaphilosophy, however, is that it is not clear how many epistemologists are signed on to it, and may not appreciate the idea that these more pragmatic considerations should be brought in to determining our best theory of knowledge. But there are other possible motivations for appealing to holistic criteria as a matter of theory selection, even among philosophers who do not adopt that metaphilosophical framework. One can find several leading philosophers appealing to holistic considerations about simplicity and the like, not in terms of the pragmatic demands of explication but rather as what they take to be a matter of sound epistemology, in following the lead of good scientific inference (e.g., Paul (2012); Nolan (2015)). Such philosophers would advocate considering simplicity/closeness-of-fit trade-offs on more general epistemological grounds, rather than metaphilosophical ones -- considerations about getting closer to the truth, rather than what we might want from our theories beyond just being true.

These philosophers all present considerations that might serve to motivate re-

opening the question of whether JTB is to be preferred as a theory of knowledge, even while accepting that the Gettier cases provide data against it. But I am not sure that the very general invocation of such holistic appeals can take us very far towards *answering* that question. I fear we will end up falling back not so much on unmootable intuitions about case verdicts, but more on ones about which of a set of theories is more elegant, and so on. Similar to the difficulties with applying Weatherson's framework, even comparing rival theories with regard to those criteria themselves may be hard enough, and the difficulty increases drastically once we look to evaluate proposed trade-offs between theories of varying elegance and similarly varying closeness to usage.

(Even if we cannot answer such questions well with Olson's Carnapian framework, it seems to me that that framework suggests instead that we should maybe consider *changing* the question. There is no need for a unique correct answer, but rather we could have a set of different knowledge-flavor reconstructions available on hand, where some might be better suited to some intellectual tasks, and others to others. I will take this point up again towards the end.)

The view I am promoting here is consistent with the above approaches, while having a distinct methodological motivation that, I hope, offers a bit more machinery to help determine how such trade-offs might best be made. In some sense, my view applies considerations of simplicity at an earlier stage of inquiry. I am drawing not on the general criteria of what makes a good theory, over and above whatever the data may tell us, but rather on a concern about figuring out just what the data are really telling us in the first place. Once one recognizes that one's data stream is itself fairly noisy, it becomes pressing to worry about the problem of *overfitting* in one's inferences and theory-selection. Some of the patterning in one's data are tracking the underlying phenomenon one wishes to capture in one's model -- but some of it is instead just tracking the noise. A tension arises between striving to capture all of the information about the real target structure on the one hand, while avoiding having one's model become itself captured by spurious twists and turns in the data.

I appealed in a recent paper (2015a) to the economist Robin Hanson's discussion of the methodological implications of having a noisy source of data. He was specifically discussing these implications in the context of ethics, and in terms of intuitions, but the point is easily seen *mutatis mutatis* to apply to other sorts of verdict data in application to other domains:

The larger one expects errors to be, the more one tends to prefer the simpler of two curves. This is because larger errors will tend to produce larger local fluctuations in data points, and these make it harder to discern local changes in the underlying curve.

In ethical "curve-fitting," one's "data" is a set of moral intuitions about what the right actions are in various particular circumstances. Regarding ethical choices made by a group, this data might consist of

intuitions from all group members, while for choices made by an individual, the data might be limited to that person's intuitions. One's "curves" are sets of ethical "principles," generally conceived. These can be very general principles, so-called "mid-level" principles, or perhaps the set of ethical choices made in certain prototypical cases (together with the relative salience of considerations used to interpolate between these cases). Together, a set of ethical principles should suggest right actions, and how they vary across some relevant range of circumstances. (Hanson (2002), 156)

More generally, to the greater extent that one expects noise in one's verdict data about some philosophical domain, to a similar extent one should prefer a simpler "curve" in terms of the structure of the philosophical generalizations that data is taken to support. As with the philosophers canvassed above, we are set up to consider a fit/simplicity trade-off, but in addition to drawing on a different set of motivations for how to think about that trade-off, we now can see a further key difference in how these considerations may be applied in practice. For the methodological considerations about curves and over-fitting offers some guidance about at least one factor that can influence how that tension between simplicity and fit should be managed. It tells us that the more we know about the nature of our source of evidence, the more clarity we can expect to achieve in gauging how best to make the trade-off. In motto form: *more noise means simpler curves*.

When I discussed these general methodological ideas in that (2015a), I raised the question, more or less in passing, of whether maybe we should take seriously the possibility that knowledge is justified true belief after all. I now want to push that point more forcefully, and not hypothetically: K=JTB really does need to be reckoned with as a still-live hypothesis. However, I don't think we are yet at a point methodologically to settle the question one way or the other. So I am not going to be arguing for the truth of K=JTB itself, so much as arguing for its live epistemic possibility at our current state of epistemological inquiry.

Of course, the general methodological point of that motto does not by itself put any pressure on epistemology or on K=JTBG at all: we must also have some information as to how noisy our verdict data is. This is where experimental philosophy becomes highly salient, for the last decade and a half of "negative program" χ -phi indicates a substantial amount of noise in our intuitive data across many areas in philosophy. I won't bother rehearsing it all here: order effects, ethnicity effects, personality trait effects, framing effects, font effects, and so on and on. The list of odd influences on philosophical verdicts continues to grow. (See, e.g., Buckwalter (2012) and Alexander & Weinberg (2014) for a recent discussions.) These results have most typically been mustered in service of various sorts of debunking arguments, either targeting some specific case, or our verdict-data practices on the whole. What I think has not been pointed out yet, though, are the methodological consequences for philosophical inferences based on such evidence. Even if one does not think these findings are a reason for a

radical revision of when and where we can legitimately appeal to verdict evidence, these negative program x -phi results provide a general reason to raise the bar significantly for when we should introduce new complications into our theory of in some philosophical domain, even when doing so would be required to capture some set of verdicts.

II. Is there Gettier-specific noise?

I have so far framed the methodological point very generally, not in any way Gettier-specifically, and indeed these concerns should have wide application. For example, such concerns should also give pause to proponents of other theories that would introduce complications into the shape of the classical theory of knowledge, such as contextualism or contrastivism. But the consequences of noise for the fit/simplicity trade-off can also be made in more local and targeted ways: to the extent that some particular region of one's data is especially noisy, to that extent one should be even more reticent to let patterns apparent in that region drive complications into one's theory. And the current state of play with regard to verdicts about Gettier cases is pretty noisy -- noisier, at least, than that background degree of clang and clatter afflicts our epistemic verdicts on the whole.

Now, it hasn't been quite the noise we were looking for, when Steve Stich, Shaun Nichols, and I found some preliminary evidence about 15 years ago that led us to suggest that intuitions about non-knowledge JTB might be culturally variable. Such cultural variability would have indeed presented a worrisome piece of noise in our epistemic usage, if it had only been, well, true. But recent attempts at replication have, very much to the contrary, revealed a fairly stable trend away from the attribution of knowledge in such cases. (See Kim & Yuan (2014) for some recent such work, as well as an overview of similar recent results; I will also discuss some of the most recent cross-cultural work on Gettier cases a bit below.) While disappointing -- no one *likes* having their results fail replication -- I think everyone should be glad to see experimental philosophy pursuing strong norms of replication here in our early days.

Yet, while substantial cross-ethnic variation would have been one serious sort of noise, had it afflicted Gettier cases, it is of course far from the only kind of noise that one might need to watch out for. And a growing set of other studies investigating Gettier cases of various sorts has shown them to be rather less clear as cases of non-knowledge than philosophical practice has taken them to be.

For starters, at least *some* specific kinds of cases that have been claimed to register as JTB-without-K have been broadly debunked, including "fake barn" type cases in particular. E.g., Colaço et al. (2014) and Turri et al. (2015); on the whole there is very little empirical evidence that fake barn-type cases are generally taken as exemplars of non-knowledge. These findings should of course not in and of themselves negatively impact our evaluation of the state of evidence regarding other particular styles of Gettier cases, which overall yield different

patterns of response (see Blouw et al. this volume). But they do of course both weaken the overall case against JTB, and raise further worries about armchair methods which had fairly widely taken them to be cases of non-knowledge. (Not without some exception being taken from the armchair itself, interestingly: see Gendler and Hawthorne (2005).)

Let us restrict our attention now to two of the most classic sorts of Gettier structures: unexpectedly defunct evidential sources, such as Russell's stopped clock; and the structure that Gettier himself innovated, in which there is in some sense divergence between a belief's truthmaker and the belief's justification. (To be clear, I'm just trying to pick out the families of structures of cases here, not provide a precise analysis of them.) The literature on the whole does clearly suggest the basic existence of what we might call a "Gettier effect", in which such structures produce lower levels of knowledge attribution than similar cases without such structures. And in some studies, this effect has been very distinct, with subjects' responses going all the way to "floor" -- very low, and indistinguishable from the responses made to other paradigms of non-knowledge, such as false beliefs and wild but lucky guesses. (E.g., Kim and Yuan (2014)). At the same time, other studies indicate that the impact of Gettier structures on folk epistemic attribution may be somewhat more complicated, and I will discuss some of those results now. First, though, let me emphasize that my argument does not actually require that there be Gettier-specific noise on evidence in the empirical literature, for the general noisiness of epistemic evaluations may be enough by itself to raise these worries about overfitting and simplicity. However, there do seem to be two ways in which some Gettierological x - ϕ work strengthens those worries. First, some studies have found Gettier cases with weak or even nonexistent Gettier effects. Call this the *intermittent effect problem*. Second, some studies have found evaluations to Gettier cases to be sensitive to factors that don't look like good candidates for inclusion in our theory of knowledge. Call this the *inappropriate sensitivity problem*.

In their important initial (2013a) attempt to demonstrate cultural uniformity of the Gettier effect, Jennifer Nagel and colleagues did indeed not find any differences across cultural groups in terms of the presence of that effect: there was a general trend across all groups in their sample to display the Gettier effect. But they also found that across all the groups they looked at, their Gettier cases produced a measurably less-than-maximal effect. Indeed they end up needing to hypothesize some substantial error theories to accommodate a challenging twist in their results: their observed Gettier effect was as strong as what they take to be an *inaccurate* epistemic effect of the mere mention of a skeptical counterpossibility, which similarly -- too similarly -- depressed the rate of knowledge attributions of their participants. Thus, while their results clearly indicate the existence of a Gettier effect, they also indicate that that effect is not similar in strength to those of more clear cases of non-knowledge, such as cases of false belief.

One prominent study that presents an intermittent effect problem is Starmans

and Friedman (2012), in which they found a robust Gettier effect for one class of case, that they call apparent evidence cases, but not in another, which they call authentic evidence cases. As they explain their distinction, in apparent evidence cases, the agents at no point have proper evidence for their belief, e.g., it has always been based on an unreliable source. But in authentic evidence cases, the agents at one time have good evidence for their belief, but then their circumstances change without their being aware of it, with a switch of truth-makers. This distinction has proved robust to several replications, though it is not at all clear that their preferred way of theorizing the distinction is the most epistemologically apt (see Nagel et al. 2013b, Starman and Friedman 2013, and Blouw et al. this volume).

In a study with an innovative "semantic integration" design, Powell et al. (2013) had subjects read vignettes that crucially involve being told that an agent *thinks* something at one point, but without use of the word "know" or similar terms, e.g., "Whatever the ultimate verdict would be, Dempsey thought Will was guilty." (In the vignette, Dempsey is a detective, and Will his prime suspect.) After a short delay and a distractor task, the subjects then perform a fill-in-the-blanks memory task, with the mental state verb missing from the crucial passage, which they are then supposed to supply from their memory of the text they had just recently read. In one of their cases (though I must add, not in all of them), Powell et al. report that in a case of Gettiered justified true belief, *nearly half* of their subjects falsely recalled seeing "know" at that point in the passage, indicating that they were at least unconsciously categorizing these cases as cases of knowledge. That is, they falsely recalled the sentence as "Whatever the ultimate verdict would be, Dempsey knew Will was guilty." The underlying idea is that such recall errors are produced by subjects' accessing their stored representation of the content of the passage, and using that stored representation of the content to answer the question, rather than a literal recollection of the text. Moreover, the subjects are also asked to make their own epistemic evaluation, e.g., whether Dempsey knew Will was guilty, or only thought Will was guilty. Subjects in a Gettier condition made knowledge attributions at the same rate as those in a parallel non-Gettiered JTB case, and both were markedly elevated by comparison with a case in which the agent's belief was stipulated to be false.

Reporting some preliminary findings from a large-scale project seeking out both uniformity and diversity across a broad range of cultures, Machery et al. (forthcoming-a) find a general Gettier effect across US, Brazilian, Indian and Japanese subjects, using a pair of structurally similar swapped-truthmaker Gettier cases but with different substantive topics (a patient in a hospital in one, and a co-worker on vacation in the other). However, while the trend is consistently one in which the Gettier cases are not knowledge, in a number of the cases they examine, they also find that the Gettier effect is observably weaker than the effect of false belief. Moreover, because their observed Gettier effects were often still some distance from the floor of false belief cases, the Machery et al. findings thus also left room for some degree of cross-nationality variation in the extent of the Gettier effects. Even though all 4 nationalities displayed the

overall trend of Gettier cases generally counting as not-knowledge, there seemed to be differences as to just how strong the effect was for the two types of cases they used. The US subjects did not treat the two Gettier cases differently, but the other three nationalities did, generally attributing more knowledge in the hospital case than in the diamond case. In both cases and across all 4 populations, it never rises even close to half of the subjects attributing knowledge; my point is not that this is evidence of any group-level variation in the existence of the Gettier effect itself, but rather in the strength and uniformity of that effect. Their results thus present a combination of both the intermittent effect problem, because the subjects's responses do not go to floor, and the inappropriate sensitivity problem, because the pattern of responses are at least somewhat variant with nationality.

A further report from that global investigation (Machery et al. (forthcoming-b)) indicates that in addition to those potential cross-cultural instabilities, Gettier cases display both order and framing effects. The authors report that they “can make the Gettier intuition compelling or underwhelming by presenting it in different contexts. In particular, people find the Gettier intuition less compelling when a case describing a justified, but false belief is presented before a Gettier case. Furthermore, we report a surprising framing effect: Two Gettier cases that differ only in their philosophically irrelevant narrative details elicit substantially different judgments.” The difference in narrative details in their study concern the source of the agent’s accidentally true beliefs. In one case, the belief has acquired via testimony, and in the other, it has been acquired via perception (it is a classic stopped-clock case). Their data show a Gettier effect that is markedly weaker in the perception case than in the testimony one. While the type of source for a belief is of course not epistemically irrelevant in general, nonetheless it is not a factor that has been anticipated to make a relevant difference with regard to the Gettier effect in particular. As such, their finding is yet another piece of *prima facie* evidence of inappropriate sensitivity.

One dramatic threat of the inappropriate sensitivity problem can be found in work that hybridizes the Gettier effect with Joshua Knobe’s side-effect effect, to produce an “epistemic side-effect effect” in which the moral valence of an action's consequence apparently impacts the extent to which subjects will report that an agent knew that those consequences would or wouldn't happen. In particular, if the actions brought about a negative moral consequence, subjects were more willing to say that the agent in the vignette knew that the consequence would happen, than if the actions brought about a positive moral consequence. But it's not just that the Knobe effect can manifest epistemically -- it apparently can produce a central trend in favor of the attribution of knowledge even in cases with classic Gettier structures (Buckwalter 2013; Beebe & Shea 2013; with some preliminary cross-cultural replication by Kim and Yuan (in prep)). One could defend the moral valence of consequences as being of epistemic relevance, as a kind of extreme form of pragmatic encroachment, but I suspect that most epistemologists would not wish to go that way.

One last report of an inappropriate sensitivity result. Turri et al. (2015) looked at a number of variations of Gettier-type structures, and they report a substantial effect of how closely the swapped truthmaker resembles the original truthmaker, in which less resemblance predicts a lower degree of knowledge attribution. Merely swapping the agent's pen with a nearly-identical one appears to be consistent with comparatively high rates of knowledge attribution; replacing Jones' owning a Ford with the rather distinct sort of fact of Brown's being in Barcelona, seems to produce very low rates. Future research may be needed to determine whether similarity/dissimilarity per se is the best way to theorize this difference, but in all it does seem that cases with otherwise similar Gettier structures can produce rather disparate rates of knowledge attribution, and if the culprit here does turn out to be a matter of the resemblance between initial and final truthmakers, that does not seem like something we would want to be reflected in our theory of knowledge itself.

In a nutshell: we find evidence that the "Gettier effect" is somewhat *transient*, appearing in some studies, and in some conditions in some studies, but not others; in many of the studies where it does appear, it is a diminished effect, not driving subjects to floor, or having an effect size of comparable scale to effects that are taken to be noise effects; and it shows evidence of sensitivity to irrelevant factors, such as the degree of resemblance between an original object and a swapped truthmaker, the moral valence of side-effects in the scenario, or nationality. All in all, the Gettier effect provides a weaker signal about knowledge than we epistemologists seem on the whole to have taken it to provide, and it seems to have its own special susceptibility to noise. Thus we have a sharpening of the main argument for taking JTB to be still a live hypothesis: in general, more noise means simpler curves, and not only is there a fair amount of noise in the verdict evidence in general, but even more so, there's evidence of even more noise in the particular vicinity of the Gettier cases.

I want to take a moment to be very clear about what is being contested here, and what is not. I am *not* trying to argue that we should all in all take the right descriptive account of the folk epistemology of the Gettier cases to be up for grabs, let alone, as being pro-JTB. Indeed, it is extraordinarily rare to find a version of switched-truthmaker or stopped-clock types of cases, in which even a small majority of subjects attribute knowledge. Even the original Weinberg et al. (2001) study reported their Asian sample only as having close to fifty-fifty verdicts here. At worst one sees a slight majority, in a handful of such versions; the epistemic side-effect effect cases are perhaps the most interesting outlier here. So it is definitely not going to be the case that one should take the verdict data here as actually pointing towards JTB. I am not here challenging the overall picture that verdicts about Gettier cases trend towards the "not knowledge" direction. In short, the experimental philosophy evidence really does point fairly clearly towards the existence of a Gettier effect on knowledge attributions. Rather, the point I am trying to press is that we have a fair amount of evidence at this time that, when one bounces these sorts of cases off human epistemometers, they often ricochet with more than a little spin on them, and can

sometimes carom off in odd directions. They bounce with varying degrees of force, and there seem to be a substantial range of ways in which odd, epistemically-inappropriate factors can push them around. And as a further point of clarification, I am also not arguing that this amount of noise in itself rises to the level that we can now declare Gettier case verdicts on the whole unreliable, and that for such a reason they therefore should be pruned away from our total set of case verdict evidence. Rather, I am looking to show how even while including these points in our total set of case verdict evidence, we have reason to worry about letting them shape our theory of knowledge to include machinery that will capture them. Perhaps the Gettier verdicts, while overall of the sort that Gettier reported them to be, are nonetheless a kind of epistemic outlier that we should not allow much say in determining which theory of knowledge is correct.

And the possibility that Gettier cases may be a kind of outlier shouldn't really be totally surprising to us, for reasons that bring me to my next supporting argument: Gettier cases are both *rare* and *weird*. They are real, and probably every epistemologist has noticed at least couple-three cases in their own lives that fit that structure. (I know I have, anyhow.) But such cases represent a pretty thin slice of our epistemic lives, at best. I will appeal here to the reader's own sense of how fairly uncommon they are -- I suspect that you've noticed them on a few occasions, but only a very few, and that they perhaps seemed particularly noteworthy to you on those occasions precisely because of their overall rarity. Note that the relevant sort of rarity does not concern the frequency of the occurrence of Gettier-type situations, but the frequency of epistemic evaluations of Gettier-type situations, in which the relevant aspects of the situation are recognized and even capable of being brought into the evaluation.

This last point is closely related to the weirdness of such cases. What I have in mind here is that these cases are very different from the most ordinary sorts of cases that we evaluate as knowledge or not, in terms of the particular sorts of information that we have to bring together in recognizing their structure. In the ordinary course of our epistemic evaluations of each other, we typically have information about an agent's take on some proposition in question, including whether or not they can be expected even to have a take in the first place; and we have information about how trustworthy we find their judgment, which also can (but much more often, doesn't) include specific information about what evidence they might cite on behalf of their view. And that's about it, isn't it? We do not, in contrast, tend to have a lot of information, or even any information at all, about any specific inferential pathway our evaluative targets may have taken to arrive at their beliefs. And it seems to me we only in the rarest of circumstances are in a situation to know that their belief might be true, while also being aware of a range of possible truth-makers for that belief, I suspect that these elements, although central to the structure of many Gettier cases, represent dimensions of everyday epistemic evaluation that we simply do not hardly ever find ourselves concerned with. It would be one thing if Gettier cases involved a highly unusual configuration of elements which individually were very commonly involved in our ordinary epistemic evaluations. But the key structural elements themselves do

not seem to me to play much of a role in our folk epistemologies.²

And here I would appeal again to our own sense of the phenomenon, this time *qua* teachers of epistemology -- do you not very often find it rather challenging to get your undergraduates even to notice these structures when their grades depend on it? And the Powell et al. results seem to indicate that often, unless folks are really hit squarely over the head with the Gettier structure, they just aren't particularly taken with them as failures of knowledge. The results of Turri (2013) are also salient here, as an example of a hit-and-miss Gettier effect: when simply presented with a Gettier case without any special window-dressing, his participants did not generally attribute knowledge, but were on the whole close to neutral as to the epistemic status of the case. He also, very interestingly, found that you could get participants to rate the case overall as an instance of non-knowledge if you went out of your way to emphasize its structure to them. So while the folk are not totally blind to possible epistemic consequences of Gettier structures, they do not seem at all attentive to them, especially in contrast to factors like the truth or falsity of the target's beliefs, or whether the agent has any plausible evidence for the proposition at all or is just guessing.³ The Starman & Friedman results, as well as what I suspect is a fairly common experience with undergraduates, suggests that this holds much more so for swapped-truthmaker sorts of cases than for unexpectedly-defunct-sources-of-evidence cases.⁴

So we have little reason to think that our ordinary capacity to render knowledge verdicts is especially well-tuned to the presence of Gettier structure of situations. And we thus both should be unsurprised that there is a lot of noise in the vicinity

² I feel obligated to note that the claims in these last two paragraphs are heavily empirically committed, and subject to potential refutation by the right kind of studies. I would very much welcome such work -- whether it confirmed my estimations here or not -- and I am working now on a study to try to get a bit of a handle on the rarity question.

³ Turri does put forward an interesting interpretation of his findings to suggest that some instances where we detect the presence of the Gettier effect should be weighted more heavily than some instances where we do not.

⁴ Jennifer Nagel made the argument to me recently at a conference that, given the rarity and weirdness of these cases, the fact that the folk epistemic evaluations are sensitive to Gettier structures *at all* suggests a kind of "poverty of stimulus" argument. If we didn't have some sort of innate epistemic predilection to be sensitive to such structures, how could we display even the intermittent and inappropriately sensitive Gettier effect that we do? It is an interesting argument, but ultimately I think it only shows that we must have some cognitive sensitivity to such structures *in general*, without that sensitivity being part of our distinctly *epistemic* cognition. The side-effect effect is perhaps a useful comparison: it shows up across lots of our cognitive lives, including our attributions of causation, but should for that reason perhaps be treated as widely-occurring potential source of noise across many (but not necessarily all) of the various domains in which it obtains.

of verdicts about Gettier cases, and, accordingly, should refuse to place any special weight on Gettier case verdicts in our inferences.

III. Conclusion and Methodological Upshots

To sum up:

1. In any sort of model-building project where we can expect noise to be present in our data, we have to wrestle with the problem of overfitting, and the more noise we expect, the less we should let each dip and turn in the data determine our choice of best model. In slogan form, more noise means simpler curves.
2. We do not *yet* have good tools in philosophy for figuring out how exactly to make such trade-offs. But the general level of noise we know to present in the verdict data is enough to make even seemingly radical such trade-offs live options.
3. In the specific vicinity of Gettier cases as a class, the experimental evidence at this time definitely points to the existence of a Gettier effect on knowledge attributions. However, even given the moderate level of noise that seems to afflict our epistemic verdicts in general, this effect seems (again, as of this time) to be both more intermittent, and more sensitive to epistemically-extraneous factors, than knowledge verdicts are on the whole. Moreover, the rarity and weirdness of Gettier cases as targets of epistemic evaluation further suggest, at a minimum, that our selection of curves need not owe any special fealty to those cases.
4. In all, as a matter of the epistemological theory of knowledge itself, and not just the psychological theory of the workings of folk epistemology, we should treat it as a live possibility that something more like JTB may be our best theory, despite the widespread existence of the Gettier effect. And we will likely not be able to settle this question until further methodological improvements are made, that will enable us to evaluate proposed fit/simplicity trade-offs responsibly.

So, to be emphatically clear on this last point, I do not think this really adds up to a positive case that we should, in fact, all things considered, at this time, endorse JTB over JTBG. We need to look out for what kind of trade-off to make between simplicity and closeness-of-fit, and given the general noisiness regarding verdicts about epistemic cases, and Gettier cases in particular, we should certainly not assume that we are in a position to sacrifice the simplicity of JTB for the apparently greater fit of some JTBG. But in our current state of methodological impoverishment, we absolutely cannot assume, either, that we should *not* make the trade in that direction. We should, rather, take ourselves to be currently in a state of unresolved ignorance regarding the epistemic significance of the Gettier cases as they manifest in our case verdicts. I will close by suggesting three distinct and complementary ways we should look to resolve that ignorance, moving forward.

First: we need more data! There has been a terrific explosion of Gettier-related results, and it is particularly encouraging that psychologists like Starmans and Friedman have recently wanted to get in on the game, and in a way that is deeply collaborative with philosophers. But for all that it is still early days yet, especially in terms of exploring empirically the different dimensions along which different Gettier cases might vary. The overall picture that develops may reveal an ever-increasing set of quirks -- or it also may instead resolve into a much more modest and stable set of effects, and ones that can make a better claim to being incorporated into our theories of knowledge. (Indeed it is always a danger in writing a chapter like this one, that by the time it sees print, the empirical tide may have turned against it. It's just a risk an author has to take, in this area; *dulce periculum est.*)

The expertise defense should be promising to explore in this particular area. That is, perhaps the Gettier-related flukes and fluctuations in the experimental data can be explained away in terms of deficits manifested by the folk when they produce their verdicts, while philosophers, due to training or acumen, may prove immune to such foibles. Now, on the whole the 'expertise defense' has fared rather poorly (Schulz et al. (2011); Machery (2012); Nado (2014); Tobia et al. (2013); Schwitzgebel and Cushman (2015); Buckwalter (forthcoming)). So one should make no presumption that the expert populations' verdict data will prove to be completely noiseless, when compared to the folk. Nonetheless, even if it seems highly unlikely that all of the intermittence and inappropriate sensitivity will be explained away in terms of the folk's lack of expertise, it is also still very plausible that at least some of it will prove to be amenable to such a treatment. And indeed, at least some results are highly suggestive in that regard. As noted above, Turri (2013) reports articulating Gettier cases to make their distinct aspects clearer and more salient can sharply decrease subjects' attributions of knowledge. And Pinillos et al. (2011) have shown that more reflective subjects may be less sensitive to the Knobe effect in general, and thus we could conjecture that they may be less susceptible to the Knobe-meets-Gettier phenomena as well. And of course, at least anecdotally, analytic philosophers tend to display a uniformly and robustly strong Gettier effect, taking such cases to be paradigm instances of non-knowledge. But on the whole, we really just do not know where philosophical expertise does or does not defuse these worries about noise in the verdict data. It will be a highly valuable empirical project to determine just where philosophical expertise can make such a difference, or where it fails to do so.

Second, and no less urgently if we are going to address the issue of the simplicity/fit tradeoff, we need to figure out how to adapt better quantitative methods of inference for philosophical inference. In particular, we need to put ourselves into a position to be able to apply quantitative modeling approaches more rigorously. There are excellent formal tools that empirical modelers use when trying to decide when it is or is not a good idea to complicate a model, such the Bayes Information Criterion or the Akaike Information Criterion, that provide a mathematical measure of the trade-off between adding additional

parameters to one's model, and how close a fit the model is to the observations. We are not yet anywhere near where we can think about the space of knowledge verdicts on the whole in such terms. One way to see this methodological deficit is to ask yourself, *just how far apart* are JTB from JTBG, in terms of their fit to the data? How big a divergence from the verdict data is JTB to begin with? After all, they agree on a considerably huge subset of verdicts. And of course we will want ultimately to be able to evaluate both the degree of fit and the degree of complexity for a range other sorts of rival hypotheses; e.g., how does SAFELY HELD TRUE BELIEF do in these terms, compared to those rivals? Or other candidate proposals for how to fill in the G in JTBG? For that matter, we will want to make sure that even JTB would win in a heads-up competition against the likes of such obviously even simpler theories like K=TB. (Or if it doesn't, then that would be a pretty interesting result, too!) Because we really have at present no way of answering questions like those that is even slightly non-handwaving, we have no way of rigorously evaluating different possible simplicity/fit trade-offs. Until we can get further along in that direction, we may have to treat the selection between JTB and its many rivals as an open choice, with a much wider space of live epistemic possibilities than we had reckoned.

That brings me to my third methodological suggestion for how we should proceed. Once we recognize that our space of theoretical possibilities is still rather more open than we may have thought, we need to explore these newly-reopened regions of epistemology space more fully. Some philosophers will, quite reasonably, want to contend that we should appeal to general theoretical results in epistemology to winnow down the hypothesis space, such that perhaps findings about, say, epistemic luck could pre-empt further consideration of JTB. I want to be clear that of course, wherever we have theoretical results that are arrived at in a largely verdict-independent way, they should of course be brought to bear as a separate constraint on our investigations. As I noted in the introduction, I am taking it as a given that abductive inferences from verdict data is *an* important method in epistemology, but not at all *the* method of epistemology.⁵ So I am not objecting to this general idea at all, that we should bring such considerations in. For example, we may take ourselves to have good reasons to treat knowledge as factive, above and beyond what the verdict data may look like, and this may play a role in fact in sorting our some apparent oddities in the verdict data in terms of protagonist projection and the like.

For all that, though, we need to be very careful about a kind of path-dependency here: we epistemologists have arrived at a state of consensus or near-consensus regarding many of these claims in no small part because of what Blouw et al. call the "business as usual" status of the Gettier verdict. It need not be the case that those claims depend *entirely* on those verdicts; they don't. But we have spent a half century now vigorously and creatively exploring the virtues of theories that

⁵ And, for example, I think that a version of Edward Craig's teleological methodology in his (1990) would be well worth bringing into greater contact with the method of cases. See, e.g., my (2015b).

respect that verdict data, and thus many of those virtues are fairly plain to view at this point. Yet we have just not done much to explore notions of knowledge that would be robust to things like epistemic luck, truthmaker-disconnection, and the like. Such recondite theories of knowledge will likely have verdict-independent theoretical virtues of their own -- epistemologically attractive features beyond how efficiently they may capture much of the verdict data. I would speculate, for example, that they may offer a satisfactory picture of how we can easily transition from knowing "P is highly likely" to knowing, simply, "P", as we seem to do often in our everyday cognition but which can be tricky to theorize in our theory of knowledge.

The noisiness in the vicinity of the Gettier cases also suggest another dimension in which we should explore epistemological theory-space. My discussion here has focused on the question of whether the verdict data in this neighborhood are really what epistemologists since Gettier himself have taken them to be, or whether they are misleading. Are they a signal of an epistemological truth that knowledge is more demanding than JTB -- an authentic signal, even if perhaps a weaker one than epistemologists have thought? Or instead, are they instead a bit of noise, something our epistemological theories should bypass, not incorporate? But those two options are not exhaustive. There is at least one other possible way to understand the status of the Gettier cases: while they do signal what really is an authentic epistemological insight, those facts are not ones ultimately best realized within our theory of *knowledge*. Something about these cases is pretty reliably registering with subjects, across a wide range of studies, as epistemically amiss, but those subjects at the same time neither universally nor at full strength, seem willing totally to reject them as knowledge. Believing truly in a luck-proof manner, or with sensitivity, or with safety, or while depending on no false lemmas... these are all epistemic desiderata in their own right, and should find a home in our epistemological theorizing whether or not they can or should play a role in theorizing knowledge itself. One further methodological upshot of taking the noise in our verdict data seriously, is that Alston's desideratum-based approach (2005) should itself be explored more thoroughly, and in particular, should be evaluated for extending its application from his own original target of justification, to knowledge itself.

There should be little doubt that a half century ago Gettier brought to light a truly fascinating piece of human psychology, and we now have the evidence to entitle us to take it to be a fairly widespread phenomenon, perhaps even a universal one. But determining how best to incorporate Gettier's psychological insight into our theory of knowledge remains a more open question than epistemologists hitherto have generally taken it to be, and one that likely cannot be answered well until we make some further methodological advances in philosophical model-building.

Works Cited

Alston, W. (2005). *Beyond "justification": Dimensions of epistemic evaluation*. Ithaca: Cornell University Press.

- Alexander, J., Mallon, R., and Weinberg, J. (2010) "Accentuate the Negative", *European Review of Philosophy*, 1, 297-314.
- Beebe, J. and Shea, J. (2013). Gettierized Knobe Effects. *Episteme*, 10: 219-240.
- Boyd, K., & Nagel, J. (2014). The reliability of epistemic intuitions. In Machery, E. and E. O'Neill, eds., *Current controversies in experimental philosophy*, Routledge, 109-27.
- Buckwalter, W. (2012). Non-Traditional Factors in Judgments about Knowledge. *Philosophy Compass*, 7(4), 278-289.
- Buckwalter, W. (forthcoming) "Intuition Fail: Philosophical activity and the limits of expertise." To appear in *Philosophy and Phenomenological Research*.
- Cappelen, H. (2012). *Philosophy without intuitions*. Oxford University Press.
- Colaço, D., Buckwalter, W., Stich, S., & Machery, E. (2014). *Epistemic intuitions in fake-barn thought experiments*. *Episteme*, 11: 199-212.
- Craig, E. (1990). *Knowledge and the State of Nature*. Oxford.
- Deutsch, M. (2015). *The Myth of the Intuitive: Experimental Philosophy and Philosophical Method*. MIT Press.
- Gendler, T., and Hawthorne, J. (2005). The real guide to fake barns: A catalogue of gifts for your epistemic enemies. *Philosophical Studies*, 124: 331-352.
- Hanson, R. (2002). Why health is not special: errors in evolved bioethics intuitions. *Social Philosophy and Policy*, 19(02), 153-179.
- Hetherington, S. (2012). The Gettier-illusion: Gettier-partialism and infallibilism. *Synthese*, 188, 217-230.
- Kaplan, M. (1985). It's not what you know that counts. *The Journal of Philosophy*, 82, 350-363.
- Kim, M., and Yuan, Y. (2014). No cross-cultural differences in the Gettier car case intuition: a replication study of Weinberg et al. 2001. *Episteme*, 12: 1-7.
- Kim, M. and Yuan, Y. (ms.) "Are Epistemic Intuitions Universal?"
- Lewis, D. (1983). *Philosophical Papers, Vol. I*. Oxford.
- Machery, E. (2012). Expertise and intuitions about reference. *Theoria. Revista de Teoría, Historia y Fundamentos de la Ciencia*, 27: 37-54.

Machery, E., Stich, S., Rose, D., Chatterjee, A., Karasawa, K., Struchiner, N., Sirker, S., Usui, N., and Hashimoto, T. (forthcoming-a) "Gettier across cultures." To appear in *Noûs*.

Machery, E., Stich, S., Rose, D., Chatterjee, A., Karasawa, K., Struchiner, N., Sirker, S., Usui, N., and Hashimoto, T. (forthcoming-b) "Gettier was framed!" To appear in E. McCready, M. Mizumoto, J. Stanley, & S. Stich (Eds.), *Epistemology for the rest of the world: linguistic and cultural diversity and epistemology*, Oxford: Oxford University Press.

Nado, J., (2014). "Philosophical Expertise," *Philosophy Compass*, 9, 631-641.

Nado, J. (2015). "Intuition, Philosophical Theorizing, and the Threat of Skepticism," in Fischer, E. and Collins, J., eds., *Experimental Philosophy, Rationalism, and Naturalism: Rethinking Philosophical Method*, Routledge.

Nagel, J. (2012). Intuitions and experiments: A defense of the case method in epistemology. *Philosophy and Phenomenological Research*, 85(3), 495-527.

Nagel, J., San Juan, V., & Mar, R. A. (2013a). Lay denial of knowledge for justified true beliefs. *Cognition*, 129: 652-661.

Nagel, J., Mar, R., & San Juan, V. (2013b). Authentic Gettier cases: A reply to Starman and Friedman. *Cognition*, 129: 666-669.

Nolan, D. 2015. The A Posteriori Armchair. *Australasian Journal of Philosophy*, 93: 211-231.

Olsen, E. 2015. Gettier and the method of explication: a 60 year old solution to a 50 year old problem. *Philosophical Studies*, 172: 57-72.

Paul, L. 2012. Metaphysics as modeling: the handmaiden's tale. *Philosophical Studies*, 160: 1-29.

Pinillos, N. Á., Smith, N., Nair, G. S., Marchetto, P., and Mun, C. (2011). Philosophy's new challenge: experiments and intentional action. *Mind & Language*, 26: 115-139.

Powell, D., Horne, Z., Pinillos, A., & Holyoak, K. J. (2013). Justified True Belief Triggers False Recall of "Knowing". *Proceedings of the 35th annual conference of the cognitive science society*, 1151-1156.

Sartwell, C. (1991). Knowledge is merely true belief. *American philosophical quarterly*, 28: 157-165.

Schulz, E., Cokely, E., and Feltz, A. (2011). "Persistent bias in expert judgments

- about free will and moral responsibility: A test of the expertise defense,” *Consciousness and Cognition*, 20: 1722-1731.
- Schwitzgebel, E., and Cushman, F., 2015, “Philosophers’ biased judgments persist despite training, expertise and reflection,” *Cognition*, 141: 127-137.
- Starmans, C., and Friedman, O. (2012). The folk conception of knowledge. *Cognition*, 124: 272-283.
- Starmans, C., and Friedman, O. (2013). Taking ‘know’ for an answer: A reply to Nagel, San Juan, and Mar. *Cognition*, 129: 662-665.
- Swain, S., Alexander, J., and Weinberg J. (2008) “The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp”. *Philosophy and Phenomenological Research*, 76, 138 – 155.
- Tobia, Kevin, Wesley Buckwalter, and Stephen Stich. (2013). "Moral intuitions: Are philosophers experts?." *Philosophical Psychology*, 26: 629-638.
- Turri, J. (2012). Is knowledge justified true belief?. *Synthese*, 184(3), 247-259.
- Turri, J. (2013). A conspicuous art: Putting Gettier to the test. *Philosopher's Imprint*, 13: 1-16.
- Turri, J., Buckwalter, W., & Blouw, P. (2015). Knowledge and luck. *Psychonomic bulletin & review*, 22, 378-390.
- Weatherson, B. (2003). What good are counterexamples? *Philosophical Studies*, 115: 1-31.
- Weatherson, B. (2014). Centrality and marginalisation. *Philosophical Studies*, 171(3), 517-533.
- Weinberg, J. (2007) “How to Challenge Intuitions Empirically Without Risking Skepticism.” *Midwest Studies in Philosophy*, 31(1), 318 – 343.
- Weinberg, J., Nichols, S. and Stich, S. (2001) *Normativity and Epistemic Intuitions*. *Philosophical Topics*, 29, 429 - 460.
- Weinberg, J. and Crowley, S. (2009) “Loose Constitutivity and Armchair Philosophy”, *Studia Philosophica Estonica*, 2.2, 177-195.
- Weinberg, J., (2015a). “Humans as Instruments; Or, The Inevitability of Experimental Philosophy”, in Fischer, E. and Collins, J., eds., *Experimental Philosophy, Rationalism, and Naturalism: Rethinking Philosophical Method*, Routledge.

Weinberg, J. (2015b). "Regress-Stopping and Disagreement for Epistemic Neopragmatists." In Henderson, D. and Greco, J., eds., *Epistemic Evaluation: Purposeful Epistemology*. Oxford.

Weinberg, J. (forthcoming). "Beyond Positive & Negative: Towards a Unified Account of Experimental Methods & Philosophical Progress". To appear in Nado, J., ed., *Advances in Experimental Philosophy and Philosophical Methodology*, Bloomsbury.

Williamson, T. (2008). *The philosophy of philosophy*. Wiley.