

# Restrictionism & Reflection: Challenge Deflected, or Simply Redirected?<sup>1</sup>

Jonathan M. Weinberg  
University of Arizona

Joshua Alexander  
Siena College

Chad Gonnerman  
Indiana University

Shane Reuter  
Georgia State University

## 0. Abstract

It has become increasingly popular to respond to experimental philosophy by suggesting that experimental philosophers haven't been studying the right kind of thing. One version of this kind of response, which we call the *reflection defense*, involves suggesting both that philosophers are interested only in intuitions that are the product of careful reflection on the details of hypothetical cases and the key concepts involved in those cases, and that these kinds of philosophical intuitions haven't yet been (and possibly cannot be) adequately studied by experimental philosophers. Of course, as a defensive move, this works only if reflective intuitions are immune from the kinds of problematic effects that form the basis of recent experimental challenges to philosophy's intuition-deploying practices. If they are not immune (or at least sufficiently less vulnerable) to these kinds of effects, then the fact that experimental philosophers have not had the right kind of thing in their sights would provide little comfort to folks invested in philosophy's intuition-deploying practices. Here we provide reasons to worry that even reflective intuitions can display sensitivity to the same kinds of problematic effects, although possibly in slightly different ways. As it turns out, being reflective might sometimes just mean being wrong in a different way.

## 1. Reflections on the Restrictionist Challenge

Here's one rather popular way of thinking about the role that philosophical intuitions play in contemporary analytic philosophy. We advance philosophical theories on the basis of their ability to explain our intuitions, defend their truth on the basis of their overall agreement with our intuitions, and justify our philosophical beliefs on the basis of their accordance with our intuitions. Intuitions are part of *standard justificatory procedure* (Bealer 1998), and are part of what defines uniquely *philosophical methodology* (Levin 2004, Goldman 2007). Let's call this the *orthodox way* of thinking about philosophical methodology.

---

<sup>1</sup> Each author contributed fully to this project. We would like to thank James Beebe, Edward Cox, the members of Indiana University's Experimental Epistemology Laboratory, and audiences at the 2008 Midsouth Philosophy Conference, the 2008 Meeting of the Society for Philosophy and Psychology and the 2011 Meeting of the Central Division of the American Philosophical Association for helpful feedback on earlier versions of this paper.

Over the years, the orthodox way of thinking about philosophical methodology has been repeatedly challenged.<sup>2</sup> Here's one recent challenge – which has become known as the *restrictionist challenge* – in a nutshell.<sup>3</sup> Our intuition-deploying practices are in good standing to the extent to which philosophical intuitions are *appropriately* sensitive – sensitive only to those things that are relevant (or, at least, that philosophers commonly take to be relevant) to the truth or falsity of the claims for which the evidence is supposed to provide evidence. Recent empirical work, however, suggests that at least some philosophical intuitions are *problematically* sensitive – sensitive to things (e.g., ethnicity, gender, affectivity, and presentation order) that haven't traditionally been thought to be relevant to the truth or falsity of philosophical claims (see, for example, Weinberg et al. 2001, Buckwalter & Stich 2011, Nichols & Knobe 2007, and Swain et al. 2008). The situation is made only worse by the fact that we currently lack the means either to correct for this kind of problematic sensitivity or even to predict from the armchair when or where else it will appear. We find ourselves in the unenviable, and ultimately untenable, epistemic position of suspecting that *some* intuitional evidence is problematically sensitive without being able reliably to predict *what* intuitional evidence is problematically sensitive. What makes this position so challenging is that almost any way of responding to it that begins by accepting the empirical results themselves seems to involve a radical departure from the orthodox way of thinking about philosophical methodology. At the very least, it seems like we must either explain away our concerns about these kinds of intuitional sensitivities or place limits on *what* intuitional evidence we employ and *when* we can safely employ it. But, placing even these sorts of *local restrictions* on our intuition-deploying practices would seem to require a *global* shift in how we think about and approach these practices. Restrictions will work only if we understand when and where they should be employed, and here it seems like our intuition-deploying practices are going to need to be informed by work in psychology, cognitive science, and the empirically-informed philosophy of mind.

It is not surprising, then, that philosophers invested in the orthodox way of thinking about philosophical methodology have found it more attractive to respond to the restrictionist challenge by either contesting the empirical results themselves or attempting to minimize their relevance.<sup>4</sup> One way to minimize the relevance is to suggest that the sensitivities are not actually problematic; e.g., if some different groups have different intuitions about knowledge, but this is due to their also having different concepts of knowledge, then such variation may well be epistemically harmless.<sup>5</sup> Another family of approaches here is to contend that, although the experimental subjects may have demonstrated unwanted sensitivities, nonetheless real philosophers doing serious philosophy in professional settings will not suffer from the same.<sup>6</sup> Here we want to consider one particular line of response of this latter kind, which we will call the *reflection defense*. Antti Kauppinen provides the most well-developed version of this

---

<sup>2</sup> See, for example, Devitt (1994), Elgin (1996), Cummins (1998), and Kornblith (1998).

<sup>3</sup> For more detailed treatments of the restrictionist challenge, see Alexander & Weinberg (2007), Weinberg (2007), and Alexander (forthcoming).

<sup>4</sup> See, for example, Hales (2006), Ludwig (2007), Williamson (2007), Cullen (2010), and Sytsma & Livengood (2010), and for discussion Weinberg et al. (2010) and Weinberg & Alexander (forthcoming). Another way of responding to the restrictionist challenge would be to argue that it is skeptical and/or self-undermining; for discussion, see Weinberg (2007).

<sup>5</sup> See, for example, Goldman (2007), Jackson (1998), Lycan (2006), Sosa (2009).

<sup>6</sup> See, for example, Hales (2006), Ludwig (2007), and Williamson (2007), and for discussion Weinberg et al. (2010) and Williamson (2011).

defense. Here's how Kauppinen (2007, 97; italics original) summarizes the defense, discussing the role that philosophical intuitions play in epistemology:

when philosophers claim that according to our intuitions, Gettier cases are not knowledge, they are not presenting a hypothesis about gut reactions to counterfactual scenarios but, more narrowly, staking a claim of how competent and careful users of the ordinary concept of knowledge would pre-theoretically classify the case in suitable conditions. The claim, then, is not about what I will call *surface intuitions* but about *robust intuitions*.

The idea behind the reflection defense is basically that philosophers are interested in *robust* philosophical intuitions – intuitions that are the product of careful reflection and consideration of the details of the hypothetical cases and the key concepts involved in those cases.<sup>7</sup> In contrast to *robust* philosophical intuitions, the philosophical intuitions which have typically been studied by experimental philosophers consist in quick, surface-level intuitions that are (it is claimed) quite unlike their robust cousins. In short, the kinds of philosophical intuitions that have been studied by experimental philosophers simply aren't relevant to philosophy's actual intuition-deploying practices and, therefore, the concerns raised about those kinds of philosophical intuitions leave philosophy's actual practices untouched and unscathed. In this way, the reflection defense aims to insulate philosophy's intuition-deploying practices from the kinds of methodological worries that form the core of the restrictionist challenge.

It is important to see that the reflection defense turns on a two-part empirical claim, which we will call the *immunity hypothesis*<sup>8</sup>: first, that the kinds of philosophical intuitions that are relevant to philosophy's intuition-deploying practices are robust in the relevant sense; and second, that robust intuitions in that sense will be sufficiently shielded from the same kinds of problematic effects that form the basis of the restrictionist challenge. Our aim here is to provide some preliminary evidence that will contribute to both rebutting and undercutting defeat for those hypotheses, by reporting some recent findings that suggest that at least some of the problematic effects do not disappear with greater reflection. It is obvious how such findings would be a source of rebutting defeat for the immunity hypothesis, since if the picture suggested by these preliminary findings turns out to be true on the whole, then the immunity hypothesis would thereby be falsified. But it is important to see how these findings can also serve as an undercutting defeater for this key commitment of the reflection defense. For such defenders of the armchair have generally taken it as simply obvious that greater reflection should make such intuitive foibles disappear. At the very least, they have clearly not felt obliged to provide any substantive empirical evidence for this commitment. And there is clearly *some* truth to the folk theory of reflection that they are tacitly endorsing, in that often, for particular cognitive tasks, greater reflection does lead to increased performance in a number of ways. But we think that results like the ones we are reporting here reveal just how insufficient that folk theory is for the task at hand. The claim that reflection *often* helps in *some* ways with *some* tasks, simply fails to

---

<sup>7</sup> The gloss on the reflection defense offered here is our own synthesis of those offered by these various authors, and should not be taken as a direct interpretation of any one of them. For example, Kauppinen does have further conditions on “robust intuitions”, but we focus here on just the reflective component.

<sup>8</sup> This is really an instance of a more general immunity hypothesis schema. For additional discussion, see Weinberg & Alexander (forthcoming).

even make likely the further claim that reflection will help address *these* restrictionist concerns about *these* intuitive tasks. The proponents of the reflection defense take themselves to be entitled to a claim that is much stronger than they have any such easy entitlement to, and the surprisingness of our results will help illustrate just how sparse their actual entitlement to the immunity hypothesis really is.

## 2. The Instability of Truetemp Intuitions

A quick review of the previous study will be helpful. Keith Lehrer (1990) challenged early versions of reliabilism by means of the now famous “Truetemp case” – a case that involves someone whose brain has been recently equipped with a reliable thermometer and computational device capable of generating temperature beliefs. Although Truetemp’s temperature beliefs are caused by a reliable cognitive process, (according to Lehrer, and agreed to by most epistemologists) it doesn’t *seem* like he knows the current temperature. In fact, it is supposed to be clear when we consider the case that Truetemp does *not* know the current temperature, and this philosophical intuition is supposed to put pressure on the idea that a person’s true belief that *p* counts as knowledge if it is caused by a reliable cognitive process. Since we have the philosophical intuition that Truetemp doesn’t know the current temperature even though he has the true belief about the current temperature, then it takes more than merely being caused by a reliable cognitive process for a true belief to count as knowledge.

But matters turn out not to be this straightforward. Swain, Alexander, and Weinberg found that subjects tend to have different epistemic intuitions about a version of the Truetemp case depending on whether they had been asked to first consider other hypothetical cases and, if so, what other hypothetical cases they had first been asked to consider. They asked people to consider the following compressed version of Lehrer’s case, which we will simply call *Truetemp*:

One day Charles was knocked out by a falling rock; as a result his brain was “rewired” so that he is always right whenever he estimates the temperature where he is. Charles is unaware that his brain has been altered in this way. A few weeks later, this brain rewiring leads him to believe that it is 71 degrees in his room. Apart from his estimation, he has no other reasons to think that it is 71 degrees. In fact, it is 71 degrees.

Some people were asked to consider *Truetemp* before evaluating any other hypothetical case. Others were asked to evaluate *Truetemp* after considering a clear case of knowledge, which we will call *Chemist*:

Karen is a distinguished professor of chemistry. This morning, she read an article in a leading scientific journal that mixing two common floor disinfectants, Cleano Plus and Washaway, will create a poisonous gas that is deadly to humans. In fact, the article is correct: mixing the two products does create a poisonous gas. At noon, Karen sees a janitor mixing Cleano Plus and Washaway and yells to him, “Get away! Mixing those two products creates a poisonous gas!”

And, still others were asked to evaluate *Truetemp* after being asked to evaluate a clear case of non-knowledge, which we will call *Coinflip*:

Dave likes to play a game with flipping a coin. He sometimes gets a “special feeling” that the next flip will come out heads. When he gets this “special feeling”, he is right about half the time, and wrong about half the time. Just before the next flip, Dave gets that “special feeling”, and the feeling leads him to believe that the coin will land heads. He flips the coin, and it does land heads.

Swain, Alexander, and Weinberg found that, when compared with people who were asked to evaluate *Truetemp* before evaluating any other cases, people who were asked to evaluate *Truetemp* after first being asked to evaluate *Chemist* were less willing to attribute knowledge in *Truetemp*, and people who were asked to evaluate *Truetemp* after first being asked to evaluate *Coinflip* were more willing to attribute knowledge in *Truetemp*. This suggests that intuitions regarding *Truetemp* can be influenced by the context in which the vignette is considered.<sup>9</sup>

These results suggest that at least some of our philosophical intuitions are sometimes sensitive to presentation order.<sup>10</sup> Proponents of the reflection defense, however, would reply first that the results are evidence at best that philosophical intuitions *that don't manifest attention to detail and careful consideration of Truetemp* are subject to presentation order effects. And they would claim, second, that philosophical intuitions that *do* manifest attention to detail and careful consideration of *Truetemp* case *won't* display this kind of problematic intuitional sensitivity, or at least only display it to a much weaker extent.

### 3. Philosophical Intuitions and Need for Cognition

Since it is that claim of appropriate insensitivity for reflective intuitions that is of crucial importance to the success of the reflection defense, let's focus on that claim. One of the difficulties we face when trying to determine whether or not this claim is true is to determine just how best to operationalize the difference between *robust* and *surface* intuitions. We need look no further than our own intellectual lives to realize that there are times when we give a great deal of care and attention to those cognitive tasks in which we are engaging and times when we don't give as much care or attention (if any at all). But, how should we define this difference so that it can be measured or expressed quantitatively?

One possibility is suggested by John Cacioppo and Richard Petty (1982). Cacioppo and Petty introduced a measure, called “need for cognition”, that they thought corresponded to the likelihood that a person would seek out, engage in, and enjoy complex and challenging cognitive tasks – that is, that they thought would correlate to a person's *intrinsic* motivation to give a great deal of care and attention to cognitive tasks that they are presented with.<sup>11</sup>

---

<sup>9</sup> See Wright (2010) for a replication, and interesting discussion, of these results.

<sup>10</sup> It is worth noting that there are other works suggesting that philosophical intuitions can be sensitive to presentation orders (e.g., Petrinovich & O'Neil 1996, Schwitzgebel & Cushman 2011, Liao et al. forthcoming), and such effects are a mainstay of a number of psychology literatures, such as persuasion, belief updating, and categorization.

<sup>11</sup> A person's need for cognition is determined on the basis of her responses to a survey with 18 self-report items like “I find satisfaction in deliberating hard and for long hours” or “I like tasks that require little thought once I've

This construct has since proved highly fruitful in psychological research, which has in turn produced a number of reasons why we might expect that a person's need for cognition would correspond to their providing robust or surface intuitions when presented with, and asked to evaluate, a given hypothetical case. First, since a person's need for cognition is supposed to represent her intrinsic motivation to engage in effortful thinking, she will find such an activity rewarding even if no external rewards are offered – she will give care and effort to the cognitive task because she experiences enjoyment and excitement from engaging in such tasks.<sup>12</sup> Second, there is evidence that a person's need for cognition is correlated with the likelihood that the person will remember the information that is provided to them and form judgments about an argument or piece of information based on the quality of the argument or the information.<sup>13</sup> And, third, people who have a high need for cognition have been shown to perform better on a variety of cognitive tasks (e.g., the American College Test, providing anagrams, and college coursework).<sup>14</sup>

Provided all of this, we decided to use subjects' need for cognition (NFC) to operationalize robust and surface intuitions. In particular, robust intuitions were operationalized as those given by subjects with relatively high NFC scores. These are individuals who are more likely to have intuitions that manifest close attention to detail and careful consideration of the relevant content of hypothetical cases that they are presented with. And surface intuitions were operationalized as those given by subjects with lower NFC scores. These are individuals who are more likely to have intuitions that don't manifest close attention to detail or careful consideration of the relevant content of hypothetical cases. Proponents of the reflection defense will thus make the prediction that higher-NFC subjects will be either less sensitive or totally insensitive to presentation order; we, in the spirit of productive contrarianism, made the opposite prediction, namely, that higher-NFC subjects would show the same kinds of intuitional sensitivity to presentation order. The surprising result: both predictions were wrong.

#### 4. The NFC study

Our subjects consisted of 167 students (82 male, 82 female, 3 unreported) attending undergraduate classes at a large Midwestern university. Their mean age was 20.93 years ( $SD = 1.66$  years).

---

learned them". The 18-statement survey is a shortened version of the survey originally used by Cacioppo and Petty; however, studies have shown that the shortened version is just as accurate a measure as the longer version of a person's need for cognition. See Cacioppo, Petty, and Kao (1984).

<sup>12</sup> See, e.g., Thompson, Chaiken, & Hazelwood (1993). Thompson and his colleagues found that people with a high need for cognition actually give less effort and care to cognitive tasks (in this case brainstorming) when they were presented with external rewards than when simply asked to engage in the cognitive task for its own sake. By contrast, people with a low need for cognition gave more effort and care to cognitive tasks when they were presented with external rewards than when they were asked to engage in the cognitive task for its own sake. These data support the claim that a person's need for cognition represents her intrinsic motivation to engage in effortful cognitive tasks.

<sup>13</sup> See, e.g., Cacioppo et al. (1983), Cacioppo et al. (1996), Chaiken (1987), Priester & Petty (1995), and Smith & Petty (1996). In particular, while people with a high need for cognition are more likely to form judgments about an argument or piece of information based on its quality, people with a low need for cognition are more likely to be influenced by external factors, peripheral cues and cognitive heuristics.

<sup>14</sup> See Petty & Jarvis (1996), Baugh & Mason (1986), Sadowski & Gulgoz (1992), and Leone & Dalton (1988).

Each subject was asked to complete a two-page survey. The first page included instructions and four hypothetical cases: *Truetemp*, *Chemist*, and *Coinflip*, as well as a version of Carl Ginet's "fake barn" case, which we'll simply call *Fakebarn*. Subjects were randomly assigned to two survey versions, which differed only in the order in which the four cases were presented, and which the Swain et al. study had indicated were the presentation orders that would yield the most dramatic intuitional instability. In one condition, subjects received this ordering: *Chemist*, *Truetemp*, *Coinflip*, and *Fakebarn*. In the other, they received the reverse: *Fakebarn*, *Coinflip*, *Truetemp*, and *Chemist*. For each case, subjects were asked to indicate the extent to which they agreed or disagreed with a target statement attributing knowledge to the protagonist using a 5-point Likert scale (1 = strongly disagree; 2 = disagree; 3 = neutral; 4 = agree; 5 = strongly agree). On the second page, subjects were asked to complete the NFC survey as well as answer several basic demographic questions (e.g., age, gender, number of philosophy courses taken, and scores on standardized college admissions exams).

The mean NFC score for our subjects was 54.71 ( $SD = 12.32$ ). We defined NFC groupings by 1 standard deviation above and below the mean; NFC scores ranging from 18-42 were designated as "low NFC" ( $N = 28$ ), those with NFC scores ranging from 43-66 were designated as "mid NFC" ( $N = 104$ ), and those with NFC scores ranging from 67-90 were designated as "high NFC" ( $N = 28$ ). (Seven subjects failed to complete the NFC survey and were thus excluded from further analysis.) According to the reflection defense, therefore, we should find that, in response to *Truetemp*, the intuitions of our low- and mid-NFC subjects will be sensitive to the case that precedes it and that the intuitions of our high-NFC subjects won't be, or at least will be to a lesser extent.

What we found is that the pattern of responses of our low- and mid-NFC subjects nicely resembles the pattern of results reported in Swain et al. (2008). Those who considered *Truetemp* after the clear case of knowledge ( $N = 64$ ,  $M = 2.53$ ,  $SD = 1.07$ ) looked somewhat less likely to attribute knowledge to Charles than those considered *Truetemp* after the clear case of non-knowledge ( $N = 68$ ,  $M = 2.88$ ,  $SD = 1.24$ ). Since the data appeared non-normal, and the results of the Shapiro-Wilk test as well as skewness and kurtosis values indicated much the same, we used a Mann-Whitney U test to determine whether this difference, and that reported below, is statistically significant. Here the difference was marginally significant:  $U = 1824.0$ ,  $.05 < p < .10$ . The Glass rank biserial correlation,  $r_g = .16$ , which is a "small" effect in Cohen's (1988) classification. These numbers suggest that the *Truetemp* intuitions of our lower-NFC subjects are trending towards a pattern like that reported in Swain et al. (2008).<sup>15</sup>

Looking at the high-NFC subjects, our results did *not* show the same sort of order effect found in the Swain et al. study, and that we found some evidence of in our low- and mid-NFC subjects. However, these results nonetheless fail to offer support for the reflection defense, either, as the high-NFC subjects' willingness to attribute knowledge to Charles did indeed tend to vary depending on the context in which the case was considered – but *in the opposite direction* of the variation trending among the responses of the lower-NFC subjects! High-NFC subjects who considered *Truetemp* after the clear case of knowledge were *more* likely ( $N = 18$ ,  $M = 2.78$ ,  $SD =$

---

<sup>15</sup> We believe that the difference in significance is due to the fact that our sample size is smaller than the Swain et al. (2008) study and that the nonparametric test we used, the Mann-Whitney U, is significantly weaker than the corresponding parametric test, the independent *t*-test.

1.22) than those who considered *Truetemp* after the clear case of non-knowledge ( $N = 10$ ,  $M = 1.80$ ,  $SD = 0.79$ ) to attribute knowledge to Charles. A Mann-Whitney U test showed this difference to be significant:  $U = 48.0$ ;  $p < .05$ ,  $r_g = -.47$ , a “medium” effect in Cohen’s classification. We note that most of the flip in the direction of instability is driven by the high-NFC subjects’ responses to the case where they saw the clear case of non-knowledge first (see section 6).<sup>16,17</sup>

The proponent of the reflection defense might try to object that we have not, after all, shown that our high-NFC subjects actually considered the cases *really, really* reflectively, and so it *might* still be the case that the instability would fade if only they were to cross some sufficiently high threshold of reflectivity. But this line of response misconstrues the dialectical situation. The argument here is an empirical one, and thus trades on probabilities, and not mere possibilities. The claim that a bit more reflection could make unwanted order effects go away was a claim that, in advance of our findings, had a modicum of empirical plausibility. The claim that a *lot* of reflection can dissolve the effects is a claim that, in the face of our findings, is an interesting hypothesis, but nothing more than that. The reflection defense will need its own empirically ascertained contrary findings in order to be an argument in good standing.<sup>18</sup>

It should also be kept in mind that this result is surprising from the armchair. We do not think that any philosopher, in the absence of having read up closely on things like contrast effect, would even have hypothesized this reverse-order effect.<sup>19</sup> As noted, we did not even initially hypothesize it ourselves. This underscores an important element of the restrictionist challenge. We just do not currently possess a very good sense of just where and how our intuitions about various sorts of cases may be susceptible to various sorts of problematic sensitivities. This is a key disanalogy between intuitions, on the one hand, and sense perception, on the other. As part of both our biological and social inheritance, we generally have a decent sense of where sense perception may be led astray. It is far from perfect, as recent work on, e.g., eyewitness testimony

---

<sup>16</sup> One interesting feature of this pattern of results is that, when we collapse all of the NFC groups together, the results don’t look much like those found in the Swain et al. (2008) study. We believe that this reflects a difference in the subject pools used for each study. Where our subjects were drawn in almost equal numbers from introductory-level and upper-level college courses, the Swain et al. (2008) study drew its subjects almost exclusively from introductory-level college courses. Since there is a reliable correlation between NFC scores and education level (Cacioppo et al. 1996, p. 217), it is likely that the SAW subjects were mostly low- to mid-NFC subjects. Since low/mid-NFC subjects and high-NFC subjects display order effects in *opposite directions*, this would help explain why it is the case that, when we collapse all of the NFC groups together in our study, our results don’t look like those found in the Swain et al. (2008) study.

<sup>17</sup> It is worth noting that we also found various gender effects and differences. To pick just two examples, a Mann-Whitney U test found that women ( $N = 82$ ;  $M = 1.85$ ,  $SD = .85$ ) were more inclined to attribute knowledge to Dave than men ( $N = 82$ ;  $M = 1.62$ ,  $SD = .84$ ) ( $U = 2781.0$ ;  $p < .04$ ;  $r_g = .17$ ) and that women, but not men ( $p > .250$ ), were more inclined to attribute knowledge to Suzy in the order in which Fakebarn was the first thought-experiment presented ( $N = 38$ ,  $M = 4.21$ ,  $SD = .78$ ) than in the order in which it was the last thought-experiment presented ( $N = 44$ ,  $M = 3.38$ ,  $SD = 1.20$ ) ( $U = 505.0$ ;  $p < .001$ ,  $r_g = .40$ ). Since the various gender effects and differences that we found appear to be orthogonal to the project of assessing the reflection defense, we won’t discuss them any further in this paper.

<sup>18</sup> Those wishing to pursue this line will want to start with the finding of Simon, Fagley, & Halleran (2004) that high NFC with sufficient depth of processing, but not high NFC alone, can moderate the effects of framing.

<sup>19</sup> One particular upshot is that this sort of result deeply problematizes the idea that armchair versions of contextualism can handle all these order effects; see Swain, Alexander, & Weinberg (2008) for some discussion of the difficulties armchair contextualists would have even with just one direction of effect.

has shown. But it is pretty good nonetheless, and is more than adequate for most normal applications of sense perception. We can thus stress one way that the restrictionist challenge is specific to intuitions, and does not necessarily need to plague other putative sources of evidence: it applies only to places where we can expect problematic intuitional sensitivity to be a threat, but where we do not have the background knowledge or practices to quarantine that threat. The threat is mostly quarantined for sense perception, and not at all for the kind of hypothetical case intuitions that the restrictionists are challenging.

## 5. The Cognitive Reflection Test

It remains open for proponents of the reflection defense to challenge the claim that a person's need for cognition is a good measure of whether they will provide robust rather than surface intuitions when asked to evaluate particular hypothetical cases. As mentioned earlier, one of the difficulties we face when trying to evaluate the reflection defense is to determine how best to operationalize the difference between robust and surface intuitions. To this end, we decided to use need for cognition scores. It remains open for someone to suggest a different (and maybe better) way of operationalizing the difference between the two. For example, Chad Gonnerman, Shane Reuter, and Jonathan Weinberg (2011) used Shane Frederick's (2005) Cognitive Reflection Test (CRT) as an alternative way to operationalize the difference.<sup>20</sup>

There are a number of reasons why we might expect that a person's cognitive reflection would correspond to her having robust or surface intuitions when presented with, and asked to evaluate, a given hypothetical case. First, since each question on the CRT test tends to trigger an initial response that is mistaken but easily fixable with a bit more thought, the test is said to measure one's "ability or disposition to resist reporting the response that first comes to mind" (Frederick 2005, 35). Thus, people with high cognitive reflection are more likely to provide judgments based on patient, deliberative and conscious processing, while people with low cognitive reflection are more likely to provide rapid judgments based on automatic, non-conscious processing. Second, cognitive reflection is correlated positively and significantly with a number of other cognitive measures including the Scholastic Achievement Test, American College Test, and Wonderlic Personnel Test (Frederick 2005, Table 4). Thus, although CRT measures a conceptually distinct characteristic—cognitive reflection—these correlations suggest that the characteristics measured by CRT overlap with characteristics measured by a number of common tests of cognitive ability.

Given all of this, CRT appears to be a reasonable way to operationalize the difference between robust and surface intuitions; we expect people's cognitive reflection to be a good measure of whether they were giving either robust or surface intuitions. People with high cognitive reflection are more likely to have intuitions that reflect patient, deliberative and conscious judgment-processes; people with low cognitive reflection are more likely to have intuitions that do not reflect patient, deliberative, and conscious judgment-processes. Proponents of the reflection defense will thus predict that higher CRT subjects will exhibit a reduced or eliminated

---

<sup>20</sup> A person's CRT score is determined by her responses to a three-item questionnaire. For example: "A bat and a ball cost \$1.10 total. The bat costs \$1.00 more than the ball. How much does the ball cost?" See Pinillos *et al.* (2011) for some pioneering work in bringing the CRT into experimental philosophy.

sensitivity to truth-irrelevant factors. Gonnerman, Reuter, and Weinberg tested this prediction using an unusual manipulation for experimental philosophy: print font.<sup>21</sup>

A brief review of the study will be helpful. They collected people's responses to a vignette and probes designed to elicit intuitions about free will and the ability to choose otherwise in a causally determined universe:

Imagine that in the next century we discover all the laws of nature, and we build a supercomputer which can deduce from these laws of nature and from the current state of everything in the world exactly what will be happening in the world at any future time. It can look at everything about the way the world is and predict everything about how it will be with 100% accuracy. Suppose that such a supercomputer existed, and it looks at the state of the universe at a certain time on March 25, 2150 AD, 20 years before Jeremy Hall is born. The computer then deduces from this information and the laws of nature that Jeremy will definitely rob Fidelity Bank at 6:00 pm on January 26, 2195. As always, the supercomputer's prediction is correct; Jeremy robs Fidelity Bank at 6:00 pm on January 26, 2195.<sup>22</sup>

After reading the vignette, people were asked to indicate their level of agreement or disagreement with the statements "Jeremy could have chosen not to rob the bank" and "When Jeremy robs the bank, he acts of his own free will", using a 5-point Likert scale (reverse of the scale used in the NFC study). For some people the vignette and probes were printed in Arial font; for others the materials were printed in *Mistral* font. The CRT and NFC task followed on separate pages (both printed in Arial font for all subjects). Compared to subjects in the Arial condition ( $N = 78$ ,  $M = 2.36$ ,  $SD = 1.33$ ), those in the Mistral condition ( $N = 73$ ,  $M = 2.78$ ,  $SD = 1.30$ ) were less inclined to say that Jeremy could have chosen otherwise. A Mann-Whitney U test showed this difference to be significant:  $U = 2295.5$ ,  $p < .04$ ,  $r_g = .19$ , suggesting sensitivity to the font of the materials.<sup>23</sup> More importantly for the debate over the reflection defense, people's robust intuitions, whether operationalized using NFC or CRT, about the ability to choose otherwise were no less susceptible to print font effects than people's surface intuitions. For instance, high-

---

<sup>21</sup> Previous research has shown that, compared to easy-to-read print fonts, hard-to-read print fonts elicit predictions of greater difficulty, skill, and effort (Song & Schwarz 2008), improve memory performance (Dieman-Yauman, Oppenheimer, & Vaughan 2011), and increase choice postponement (Novemsky, Dhar, Schwarz, & Simonson 2007).

<sup>22</sup> The vignette and probes were borrowed from Nahmias, Morris, Nadelhoffer, & Turner (2005). See also Nahmias, Morris, Nadelhoffer, & Turner (2006).

<sup>23</sup> Some proponents of the reflection defense might try to find hope in these results. Their thought might go: hard-to-read fonts tend to trigger deeper reflection and with deeper reflection comes a greater likelihood of having the correction intuition, so the right way to respond to these results is that the hard-to-read font tends to put subjects in a better position to come to the correct intuition, in this case, that Jeremy couldn't have chosen otherwise. This kind of thought seems quite popular, but it's one that the proponent of the reflection defense shouldn't find much hope in. For starters, the reflection defense, as we use the expression, essentially involves the claim that robust intuitions, in comparison with surface intuitions, are insensitive (or at least considerably less sensitive) to the factors in question. But, as noted in the text, this claim doesn't appear to be true when it comes to robust intuitions regarding Jeremy's ability to choose otherwise and print fonts. Moreover, and perhaps more importantly, the hypothesis that the hard-to-read font managed to trigger deeper reflection doesn't fit too well with the results. For it were true, the high-CRT/Arial answers should differ from the low-CRT/Arial answers in at least roughly the same way that low-CRT/Mistral answers differ from low-CRT/Arial answers. But this prediction is not borne out in the data. For the details, see Gonnerman, Reuter, & Weinberg (2011).

CRT subjects (those who scored 3 out of 3 on the test) in the Mistral condition ( $N = 19$ ,  $M = 3.16$ ,  $SD = 1.38$ ) were less inclined to say that Jeremy could have chosen otherwise than those who saw the easy-to-read font ( $N = 15$ ,  $M = 2.07$ ,  $SD = 1.33$ ). A Mann-Whitney U test showed this difference to be significant as well:  $U = 78.0$ ,  $p < .03$ ,  $r_g = .45$ .

However, it might well be the case that there is a better way to understand the difference between robust and surface intuitions and it might well be that, on this way of understanding the difference, robust intuitions *aren't* as unstable as surface intuitions. So, we should continue to think of different ways of operationalizing the proposed difference and to test whether robust intuitions are less susceptible than surface intuitions to problematic sensitivity. We used NFC to operationalize this difference, while Gonnerman et al. used both NFC and CRT. Both studies revealed worrisome patterns of intuitional sensitivity among surface and robust intuitions. This is not the last word; however, it does put pressure on the proponent of the reflection defense. On one reasonable way of operationalizing the difference between robust and surface intuitions (CRT scores), the robust intuitions discussed here were found to be no less susceptible to one problematic factor (print fonts), and, on the other reasonable way of operationalizing the difference (NFC scores), they were found to be no less susceptible to two problematic factors (print fonts and presentation orders).

## 6. Two Hypotheses, and Their Mixed Restrictionist Implications

Our discussion thus far has focused largely on the presence or absence of a particular sensitivity to individual differences or context conditions, viz., sensitivity to presentation order and whether it can be made to disappear with greater reflection. But there is another piece to the restrictionist challenge to consider, namely, whether the observed sensitivity need be understood as *unwanted*, a sensitivity to some non-truth-tracking factors. One might argue that this kind of sensitivity is perfectly welcome by suggesting that the relevant facts themselves really do change along with the order of the cases considered, perhaps using contextualist machinery. We do not find this move attractive, as we have argued elsewhere.<sup>24</sup> But the question can usefully be raised as to whether some *other* explanation might be available that would be less inimical to the methodological self-understanding of proponents of the orthodox view. Which in turn raises the more general question: just what explains these results, anyway? They are, as we noted above, at least *prima facie* somewhat weird.

One possible explanation might have to do with the dynamic nature of the human categorizing system, and what might be some plausible variations in such systems across more and less reflective persons. Remember, the crux of our observed effect was the differential impact of seeing *Coinflip* before seeing *Truetemp*: in low- and mid-NFC subjects, this order seemed to produce an increase in attribution of knowledge to Charles, but in high-NFC subjects, it seemed to produce a decrease in such attributions. It may be that we are looking at a *saliency* effect here, with the categorizing systems of subjects of different NFC levels being differentially affected by seeing *Coinflip* before seeing *Truetemp*. There are at least two key epistemological characteristics of Dave the coin-flipper that would likely be salient to our subjects, but maybe not equally salient to all of them. First, Dave is forming a belief by an *unreliable* method, namely, on

---

<sup>24</sup> See discussion mentioned in n.19 above – which is further complicated by our result that there may well be different average reactions to order, depending on factors like one's need for cognition.

a “signal” that really has no informational value whatsoever. Second, Dave is forming a belief *unthinkingly*, on the basis of a “special feeling”, without any discursive reasons. So here is one hypothesis. Perhaps the first characteristic of Dave is more salient to low- and mid-NFC people, whereas the second is more salient to high NFC people. After all, high-NFC people are precisely those who appear to value effortful cognition more than lower NFC people. It is reasonable to suggest, therefore, that they are more likely to pay attention to Dave’s lack of discursive reasons for believing what he does.<sup>25</sup> For the low- and mid-NFC folks, then, Dave sets up a positive comparison for Charles: *hey, at least this thermometer guy’s belief is being formed reliably, better than that coinflipper guy’s belief!* Whereas for the high-NFC folks, Dave sets up a negative comparison for Charles: *dang, this thermometer guy’s belief is just as unreasonable as the coinflipper guy’s belief!* Hence, the increased attribution from the former folks, as compared to the level of attribution from the latter folks.

We consider this hypothesis to be basically an unfriendly one for the proponent of the orthodox view. It is hard to see how it could help unless there was reason to think that the categorizing systems of one NFC grouping generally gives rise to better, perhaps more reliable, intuitions than the systems of the other NFC groupings.<sup>26</sup> This is a possibility. In section 3, we note some of the empirical work suggesting ways in which high-NFC correlates with various epistemically desirable tendencies and accomplishments. Perhaps within these works there are also reasons to be found and articulated for preferring the intuitions of high-NFC individuals. Still it is important to note that those reasons won’t entail that we ought to accept their intuitions wholesale, no questions asked. Remember, there is also reason to think that at least some of their philosophical intuitions are problematically sensitive. Consider, again, the NFC study. It seems unlikely that whether Charles really knows depends on whether a case like *Chemist* or *Coinflip* first came up, at least not in the manner that the intuitions of our high-NFC subjects would tend to suggest. (And there is little reason to think that problematic intuitions on the part of high-NFC subjects are likely to stop there. For example, Petty and Jarvis (1996) note that there is good reason to think that high-NFC individuals are *more* likely to be susceptible to at least some priming effects than low-NFC individuals.) Even this hypothesis, then, leads us to the standard restrictionist refrain: to explore and defend such a hypothesis, we would need experimental work to document just where the problematic intuitions arise; that is, we would need to do some experimental philosophy on the (assumed) cognitive elite. After all, the intuitional susceptibilities of high-NFC individuals reported here were not observed from the armchair. Indeed, when it comes to the effects of presentation orders on knowledge intuitions, the philosophers we have seen who have hypothesized about them have tended only to offer predictions of the sort reported in Swain et al., and not the reverse-direction results reported here.

This categorizer-system variation hypothesis is not the only plausible hypothesis in this area, though. Interestingly, social psychologists working on priming effects and need for cognition have reported results similar to ours, namely, cases where high- and low-NFC subjects show inverted reactions to priming. For example, it has been reported that when primed with the

---

<sup>25</sup> The idea that there are attentional differences between high and low NFC subjects is not without precedent in the empirical literature. For example, Timothy Osberg (1987) found that NFC scores are positively correlated with scores on a measure of private self-consciousness, “the tendency to attend to and reflect on one’s own attitudes and feelings” (p. 443; see also Cacioppo et al., 1996, p. 216).

<sup>26</sup> Certainly, it is unlikely that the truth-values of the corresponding knowledge claims vary across NFC groupings.

concept of persistence, low-NFC subjects tended to form a more favorable impression of an ambiguously described individual than when primed with the concept of stubbornness while high-NFC subjects displayed the opposite pattern (Martin, Seta, and Crelia 1990). Similar results have been reported with respect to emotional states and event estimations (DeSteno, Petty, Wegner, and Rucker 2000): When first asked to recall a time in which they were very sad low-NFC subjects tended to estimate more sad events in the future (e.g., more people will have to euthanize a pet this year) than when asked to recall a very happy or angry time; again, high-NFC subjects went the other way. So finding oppositely directed variation among low- and high-NFC subjects is not without precedent. These are but two examples.

However, these two report results that are in an important way different from ours. They found a greater tendency toward *assimilation* among low-NFC subjects and toward *contrast* among high-NFC subjects. We found the reverse direction of effects. For example we found that when first presented with a clear case of non-knowledge, relatively low-NFC subjects tended to attribute knowledge to Charles, rather than deny knowledge. But even the reverse can be explained, given the conditions of our survey. One way to explain the above social-psychological results is to see the low-NFC subjects as providing the cognitively least demanding, or *default*, response and high-NFC subjects as giving the more demanding, *correction-attempted*, response. It's just that high-NFC subjects overcorrected.<sup>27</sup> If that is right, then we could expect circumstances of flipped variation, just like what we found. It would require the default effect of the prime to be contrast (rather than assimilation).

There is experimental evidence demonstrating that the default could go that way. When subjects are briefly presented the name of an extremely athletic individual (in this case, 'Michael Jordan') they tend to give lower assessments of their own athletic abilities compared to the presentation of an extremely non-athletic individual ('Pope John Paul') (Musseiler, Rüter, and Epstude 2004). Exactly what factors determine the default process looks complicated, but evidence suggests that one tendency is this: when an extreme example precedes a task that requires evaluation of a particular, the default effect is contrast.<sup>28</sup> Our experimental design presented subjects with either a clear case of knowledge or non-knowledge before *Truetemp*. The social psychological work on prime extremity suggests that these conditions make contrast the default response, especially given the *very* poor epistemic standing of coin-flipper Dave. If so, borrowing only the basics from the social psychologists' explanations, we can explain why relatively low-NFC subjects will be more inclined to give the default, contrasting response, as found. This explanation works as long as the primes are blatant enough to push the lower-NFC subjects one way and to trigger the high-NFC subjects' correction processes.

This *reflective overcorrection* hypothesis is a very interesting one from the point of view of the restrictionist challenge, because it would seem to contain a strongly armchair-friendly element. For, if the high-NFC subjects are reacting to what they perceive to be an inappropriate influence of the previous case on their judgment about a current case, then that means that at least highly reflective persons can, in real time and without any special training, become aware of the threat of such order effects, and try to take appropriate compensatory steps. And if it turns out that

---

<sup>27</sup> This is the explanation that the researchers themselves gave. See, for example, DeSteno et al. (2000, especially pages 408-409). See also Petty, DeMarree, Briñol, Horcajo, and Stathman (2008).

<sup>28</sup> For more on default contrast effects, see Musseiler, Rüter, & Epstude (2004) and Petty et al. (2008).

reflective persons have such error-detection capacities, then such capacities could contribute substantially towards philosophers having practices that could thus be sufficiently shielded from such errors, just as proponents of the orthodox view would hope.

It is, nonetheless, an *overcorrection* hypothesis. Even on this hypothesis, our subjects, at least, were not on the whole able to make a correctly calibrated compensation for the influence of the order of the vignettes. Such a finding is generally in keeping with the psychological literature on debiasing, in which merely making people aware of a given foible in reasoning can sometimes work, but also is often insufficient on its own to enable them to compensate well.<sup>29</sup> It may turn out that reflective agents do indeed have some of the right raw materials for making a shield from restrictionist worries – but that forging real protection from those materials may require painstaking scientific investigation. Proponents of the orthodox view may take some hope from this hypothesis, but at best it is hope for the future.

At this point, there are no explanations to offer that are not mere hypotheses. There are a number of different possible implications of this research, and for that matter, it is always important to keep in mind in these discussions that results like the ones reported here are themselves only preliminary. We take ourselves to have contributed one more plank to the growing empirical foundations of an account of the psychology of philosophical intuition, in a way that adds a bit of strength to the restrictionist challenge, in part by showing how nontrivial – and, at this point, how at variance with the little bit of evidence that there is – the empirical commitments of the reflection defense really are. Proponents of the orthodox view have been looking for a way to dismiss the restrictionist challenge, to make legitimate their desire to be able to ignore the empirical dimensions of traditional armchair philosophical methodology. Our results will hopefully contribute to their coming to see that this will not be possible. If the armchair is to be saved from experimental philosophy, it will take even more experimental philosophy to do it.

## REFERENCES

Alexander, Joshua Forthcoming. *Experimental Philosophy: An Introduction*, Cambridge: Polity Press.

Alexander, Joshua and Jonathan M. Weinberg 2007. “Analytic Epistemology and Experimental Philosophy,” *Philosophy Compass*, 2: 56-80.

Babcock, Linda and George Loewenstein 1997. “Explaining Bargaining Impasse: The Role of Self-Serving Biases,” *The Journal of Economic Perspectives*, 11: 109-26.

Baugh, Bernard T. and Susan E. Mason 1986. “Need for Cognition Related to Time Perception,” *Perceptual and Motor Skills*, 62: 540-542.

Bealer, George 1998. “Intuition and the Autonomy of Philosophy,” in M.R. DePaul and W. Ramsey, eds., *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, Lanham, MA: Rowman & Littlefield, 201-239.

---

<sup>29</sup> See, e.g., Fischhoff (1982), Babcock & Loewenstein (1997), and Wilson et al. (2002).

- Buckwalter, Wesley and Stephen Stich 2011. "Gender and the Philosophy Club," *The Philosophers Magazine*, 52: 60-65.
- Cacioppo, John T. and Richard E. Petty 1982. "The Need for Cognition," *Journal of Personality and Social Psychology*, 42: 116-131.
- Cacioppo, John T., Richard E. Petty, Jeffrey A. Feinstein, and Blair G. Jarvis 1996. "Dispositional Differences in Cognitive Motivation: The Life and Times of Individuals Varying in Need for Cognition," *Psychological Bulletin*, 119: 197-253.
- Cacioppo, John T., Richard E. Petty, and Chuan F. Kao 1984. "The Efficient Assessment of Need for Cognition," *Journal of Personality Assessment*, 48: 306-307.
- Cacioppo, John T., Richard E. Petty, and Katherine J. Morris 1983. "Effects of Need for Cognition on Message Evaluation, Recall, and Persuasion," *Journal of Personality and Social Psychology*, 45: 805-818.
- Chaiken, Shelly 1987. "The Heuristic Model of Persuasion," in M.P. Zanna, J.M. Olson, and C.P. Herman, eds., *Social Influence: The Ontario Symposium* (Volume 5), Hillsdale, NJ: Lawrence Erlbaum, 3-39.
- Cohen, Jacob 1988. *Statistical Power Analysis for the Behavioral Sciences*, Hillsdale, NJ: Erlbaum.
- Cullen, Simon 2010. "Survey-Driven Romanticism," *Review of Philosophy and Psychology*, 1: 275-296.
- Cummins, Robert 1998. "Reflection on Reflective Equilibrium," in M.R. DePaul and W. Ramsey, eds., *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, Lanham, MA: Rowman & Littlefield, 113-127.
- DeSteno, David, Richard E. Petty, Duane T. Wegener, and Derek D. Rucker 2000. "Beyond Valence in the Perception of Likelihood: The Role of Emotion Specificity," *Journal of Personality and Social Psychology*, 78: 397-416.
- Devitt, Michael 1994. "The Methodology of Naturalistic Semantics," *Journal of Philosophy*, 91: 545-572.
- Dieman-Yauman, Connor, Daniel M. Oppenheimer, and Erikka B. Vaughan 2011. "Fortune Favors the **Bold** (and the *Italicized*): Effects of Disfluency on Educational Outcomes," *Cognition*, 118: 111-115.
- Elgin, Catherine Z. 1996. *Considered Judgment*, Princeton: Princeton University Press.

- Fischhoff, Baruch 1982. "Debiasing", in Daniel Kahneman, Paul Slovic, & Amos Tversky (eds.), *Judgment under uncertainty: Heuristics and biases*, New York: Cambridge University Press, 422–444.
- Frederick, Shane 2005. "Cognitive Reflection and Decision Making," *Journal of Economic Perspectives*, 19: 25-42.
- Goldman, Alvin I. 2007. "Philosophical Intuitions: Their Target, Their Source, and Their Epistemic Status," *Grazer Philosophische Studien*, 74: 1-26.
- Gonnerman, C., Shane Reuter, and Jonathan M. Weinberg 2011. "More Oversensitive Intuitions: Print Fonts and Could Choose Otherwise," Paper Presented at the One Hundred Eighth Annual Meeting of the American Philosophical Association, Central Division, Minneapolis, MN.
- Hales, Steven D. 2006. *Relativism and the Foundations of Philosophy*, Cambridge, MA: MIT Press.
- Jackson, Frank 1998. *From Metaphysics to Ethics: A Defense of Conceptual Analysis*, Oxford: Clarendon Press.
- Kassin, Saul M., Marisa E. Reddy, M., and William F. Tulloch 1990. "Juror Interpretations of Ambiguous Evidence: The Need for Cognition, Presentation Order, and Persuasion," *Law and Human Behavior*, 14: 43-55.
- Kauppinen, Antti 2007. "The Rise and Fall of Experimental Philosophy," *Philosophical Explorations*, 10: 95-118.
- Kornblith, Hilary 1998. "The Role of Intuition in Philosophical Inquiry: An Account with No Unnatural Ingredients," in M.R. DePaul and W. Ramsey, eds., *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, Lanham, MA: Rowman & Littlefield, 129-141.
- Lahroodi, Reza 2007. "Evaluating Need for Cognition: A Case Study in Naturalistic Epistemic Theory," *Philosophical Psychology*, 20: 227-245.
- Lehrer, Keith 1990. *Theory of Knowledge*, Boulder, CO: Westview.
- Leone, Christopher and Charles H. Dalton 1988. "Some Effects of the Need for Cognition on Course Grades," *Perceptual and Motor Skills*, 67: 175-178.
- Levin, Janet 2004. "Evidential Status of Philosophical Intuitions," *Philosophical Studies*, 121: 193-224.
- Liao, S. Matthew, Alex Wiegmann, Joshua Alexander, and Gerard Vong manuscript. "Putting the Trolley in its Place: Experimental Philosophy and the Loop Case.

- Ludwig, Kirk 2007. "The Epistemology of Thought Experiments: First Person versus Third Person Approaches," *Midwest Studies in Philosophy*, 31: 128-159.
- Lycan, William G. 2006. "On the Gettier Problem," in S. Hetherington, ed., *Epistemology Futures*, Oxford: Oxford University Press, 148-168.
- Martin, Leonard L., John J. Seta, and Rick A. Crelia 1990. "Assimilation and Contrast as a Function of People's Willingness and Ability to Expend Effort in Forming an Impression," *Journal of Personality and Social Psychology*, 59: 27-37.
- Meyers-Levy, Joan and Laura A. Peracchio 1992. "Getting an Angle in Advertising: The Effect of Camera Angle on Product Evaluation," *Journal of Marketing Research*, 29: 454-461.
- Mueller, John H., Tim Keller, and Alison Dandoy 1989. "Face Memory and Spontaneous Personality Processing", *Bulletin of the Psychonomic Society*, 27: 103-106.
- Mussweiler, Thomas, Katja Rüter, and Kai Epstude 2004. "The Man Who Wasn't There: Subliminal Social Comparison Standards Influence Self-evaluation," *Journal of Experimental Social Psychology*, 40: 659-696.
- Nahmias, Eddy, Stephen Morris, Thomas Nadelhoffer, and Jason Turner 2005. "Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility," *Philosophical Psychology*, 18: 561-584.
- Nahmias, Eddy, Stephen Morris, Thomas Nadelhoffer, and Jason Turner 2006. "Is Incompatibilism Intuitive?" *Philosophy and Phenomenological Research*, 73: 28-53.
- Nichols, Shaun and Joshua Knobe 2007. "Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions," *Noûs*, 41: 663-685.
- Novemsky, Nathan, Ravi Dhar, Norbert Schwarz, and Itamar Simonson 2007. "Preference Fluency in Choice," *Journal of Marketing Research*, 44: 347-356.
- Osberg, Timothy M. 1987. "The Convergent and Discriminant Validity of the Need for Cognition Scale," *Journal of Personality Assessment*, 51: 441-450.
- Petty, Richard E., Kenneth G. DeMarree, Pablo Briñol, Javier Horcajo, and Alan J. Strathman 2008. "Need for Cognition Can Magnify or Attenuate Priming Effects in Social Judgment," *Personality and Social Psychology Bulletin*, 34: 900-912.
- Petty, Richard E. and Blair G. Jarvis 1996. "An Individual Difference Perspective on Assessing Cognitive Processes," in N. Schwarz and S. Sudman, eds., *Answering Questions: Methodology for Determining Cognitive and Communicative Processes in Survey Research*, San Francisco: Jossey-Bass Publishers, 221-257.

- Pinillos, N. Ángel, Nick Smith, G. Shyam Nair, Peter Marchetto, Cecilea Mun 2011. "Philosophy's New Challenge: Experiments and Intentional Action," *Mind and Language*, 26: 115-139.
- Priester, Joseph R. and Richard E Petty 1995. "Source Attributions and Persuasion: Perceived Honesty as a Determinant of Message Scrutiny," *Personality and Social Psychology Bulletin*, 21: 637-654.
- Sadowski, Cyril and Sami Gulgoz 1992. "Association of Need for Cognition and Course Performance," *Perceptual and Motor Skills*, 74: 498.
- Schwitzgebel, Eric, and Fiery Cushman forthcoming. "Expertise in Moral Reasoning? Order Effects on Moral Judgments in Professional Philosophers and Non-Philosophers," *Mind & Language*.
- Simon, Andrew F., N.S. Fagley, and Jennifer G. Halleran 2004. "Decision Framing: Moderating Effects of Individual Differences and Cognitive Processing," *Journal of Behavioral Decision Making*, 17: 77-93.
- Smith, Stephen M. and Richard E. Petty 1996. "Message Framing and Persuasion: A Message Processing Analysis," *Personality and Social Psychology Bulletin*, 22: 257-268.
- Song, Hyunjin, and Norbert Schwarz 2008. "If it's Hard to Read, it's Hard to Do: Processing Fluency Affects Effort Prediction and Motivation," *Psychological Science*, 19: 986-988.
- Sosa, Ernest 2009. "A Defense of the Use of Intuitions in Philosophy," in D. Murphy and M. Bishop, eds., *Stich and His Critics*, Oxford: Wiley, 101-112.
- Strull, Thomas, K., Meryl Lichtenstein, and Myron Rothbart 1985. "Associative Storage and Retrieval Processes in Person Memory," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11: 316-345.
- Swain, Stacey, Joshua Alexander, and Jonathan M. Weinberg 2008. "The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp," *Philosophy and Phenomenological Research*, 76: 138-155.
- Sytsma, Justin and Jonathan Livengood Forthcoming. "A New Perspective Concerning Experiments on Semantic Intuitions," *Australasian Journal of Philosophy*.
- Thompson, Erik P., Shelly Chaiken, and Douglas J. Hazelwood 1993. "Need for Cognition and Desire for Control as Moderators of Extrinsic Reward Effects: A Person Situation Approach to the Study of Intrinsic Motivation," *Journal of Personality and Social Psychology*, 64: 987-999.
- Weinberg, Jonathan M. 2007. "How to Challenge Intuitions Empirically Without Risking Skepticism," *Midwest Studies in Philosophy*, 30: 318-343.

Weinberg, Jonathan M., Chad Gonnerman, Cameron Buckner, and Joshua Alexander 2010. "Are Philosophers Expert Intuiters?" *Philosophical Psychology*, 23: 331-355.

Weinberg, Jonathan M., Shaun Nichols, and Stephen Stich 2001. "Normativity and Epistemic Intuitions," *Philosophical Topics*, 29: 429-460.

Weinberg, Jonathan M. and Joshua Alexander Forthcoming. "The Challenge of Sticking With Intuitions Through Thick and Thin," in Anthony Booth and Darrell Rowbottom, eds., *Intuition*, Oxford: Oxford University Press.

Williamson, Timothy 2007. *The Philosophy of Philosophy*, Oxford: Blackwell.

Williamson, Timothy 2011. "Philosophical Expertise and the Burden of Proof," *Metaphilosophy* 42: 215-229.

Wilson, Timothy, David Centerbar, and Nancy Brekke 2002. "Mental contamination and the debiasing problem," in Thomas Gilovich, Dale Griffin, and Daniel Kahneman (eds.), *The psychology of judgment: Heuristics and biases*. New York: Cambridge University Press, pp. 185-200.

Wright, Jennifer Cole 2010. "On Intuitional Stability: The Clear, the Strong, and the Paradigmatic," *Cognition*, 115: 419-503.