

THE PAST AND FUTURE OF EXPERIMENTAL PHILOSOPHY

Thomas Nadelhoffer and Eddy Nahmias¹

Experimental philosophy is the name for a recent movement whose participants use the methods of experimental psychology to probe the way people think about philosophical issues and then examine how the results of such studies bear on traditional philosophical debates. Given both the breadth of the research being carried out by experimental philosophers and the controversial nature of some of their central methodological assumptions, it is of no surprise that their work has recently come under attack. In this paper we respond to some criticisms of experimental philosophy that have recently been put forward by Antti Kauppinen. Unlike the critics of experimental philosophy, we do not think the fledgling movement either will or should fall before it has even had a chance to rise up to explain what it is, what it seeks to do (and not to do), and exactly how it plans to do it. Filling in some of the salient details is the main goal of the present paper.

KEYWORDS experimental philosophy; conceptual analysis; meta-philosophy; folk intuitions

1. Introduction

Experimental philosophy is the name for a recent movement whose participants use the methods of experimental psychology to probe the way people make judgments that bear on debates in philosophy. Although the movement has a name, it includes a variety of projects driven by different interests, assumptions, and goals. Just in the past few years philosophers have carried out experimental work in areas as diverse as epistemology, action theory, free will and moral responsibility, the philosophy of language, ethics, the philosophy of law, the philosophy of mind, and the philosophy of science.² All of this work shares a commitment to using controlled and systematic experiments to explore people's intuitions and conceptual usage and to examine how the results of such experiments bear on traditional philosophical debates. But, as we will explain below, while some experimental philosophers use data about ordinary intuitions to support philosophical theories, others use such data to better understand the psychological mechanisms that generate such intuitions, while still others gather such data to show that some intuitions may be too unreliable to support philosophical theories in the first place.³

Given both the breadth of the research being carried out by experimental philosophers and the controversial nature of some of their central methodological assumptions, it is no surprise that their work has recently come under attack. While we admittedly share some of the worries expressed by critics of experimental philosophy, we nevertheless

disagree with many of their reservations, including those expressed by Antti Kauppinen in this issue. We certainly do not think that experimental philosophy either will or should *fall* before its practitioners have had a chance to rise up to explain what it is, what it seeks to do (and not to do), and exactly how it plans to do it. Filling in some of these details is one of the goals of the present paper. Along the way, we will flesh out important differences between the various projects being undertaken by experimental philosophers, examine and respond to the main criticisms put forward by Kauppinen, and suggest directions for future work in experimental philosophy.

On our view, the burgeoning experimental philosophy movement has already made a significant contribution to philosophy by highlighting some important issues that philosophers need to address more fully, such as (i) **the nature and evidentiary status of both pre-philosophical intuitions and philosophically informed intuitions**, (ii) **the proper relationship between intuitions, concepts, and theories**, (iii) **the proper methods and goals of conceptual analysis**, and (iv) **the scope and limits of reflective equilibrium**. By using (philosophically) unconventional methods to bring these meta-philosophical questions to the forefront, experimental philosophers have not only reinvigorated several important debates in philosophy, they have also opened new avenues of interdisciplinary research and dialogue. Regardless of what becomes of experimental philosophy in the future, we think it has already made an important and enduring contribution to philosophy as a whole by challenging some of the basic methodological assumptions that are woven into the fabric of how many philosophers go about doing contemporary analytic philosophy.⁴

2. The Past: Three Projects of Experimental Philosophy

The critics of experimental philosophy often treat it as a monolithic project driven by the shared methodologies and goals of the experimental philosophers themselves. If this were the case, then experimental philosophy would be much easier for its critics to bring down with a single blow. But upon closer examination we find that this is decidedly not the case. So, one strategy for defending experimental philosophy against its critics in this paper will be to divide it into different projects and then examine the criticisms as they individually apply—or fail to apply—to each of these projects. Sometimes, responding to one criticism will require accepting another; conversely, criticisms that apply to one project will sometimes fail to apply to others. As we will see, the diversity of experimental philosophy is demonstrated by the fact that some criticisms require envisioning it in a way that contradicts the way other criticisms envision it. And while it may be that all of the projects within experimental philosophy will not stand together—at least not as they are presently construed—we are nevertheless confident that the entire movement will not be brought down by any of the current objections.⁵

Given the diversity of the projects carried out under the rubric of ‘experimental philosophy,’ we should state as precisely as possible both what is unique about the fledgling movement and what its participants have in common. For instance, what distinguishes experimental philosophers from not only similarly-minded naturalistic and empirically informed philosophers but also from experimental psychologists? Though the boundary here is blurry, the primary difference is that experimental philosophers actually run their own studies to get at the data they need and then show why these data are philosophically interesting.⁶ Generally, they design surveys to test laypersons’ intuitive judgments about

philosophical issues, they run the surveys, and they statistically analyze the results. But whereas the 'experimental' part of the name refers to the fact that they run studies and collect data concerning folk intuitions, the 'philosophy' part refers to the fact that they discuss the various implications these data have for philosophical debates.

Contrary to what some critics have claimed, experimental philosophers do not naively believe that their results will solve philosophical debates and 'put an end to vain quarrel' (Kauppinen 2007, 100). As we will see, their goals and ambitions are typically more modest than these critics suggest. Any areas of philosophy that rely on (i) intuition-pumps and thought experiments, (ii) appeals to commonsense and pre-philosophical intuition or (iii) conceptual analysis based in part on ordinary usage or 'platitudes'⁷ are ripe for investigation by experimental philosophers who are, above all, interested in examining these things in a controlled and systematic way. This interest is driven by two assumptions: First, there is a shared distrust of philosophers' (common) claims of the general form 'X is intuitive,' 'Ordinarily, we believe X,' 'The ordinary use of "X" is Y,' 'It is natural for people to believe X,' and the like. These are claims that philosophers usually make based upon armchair reflection on their own intuitions and (perhaps selective) consideration of their conversations with friends, family, and especially students.⁸ But these methods of determining what is widely accepted seem highly susceptible to well-known biases such as confirmation bias,⁹ disconfirmation bias,¹⁰ and false consensus bias¹¹—biases that may be even more pronounced in philosophers than laypersons given the philosophers' extensive theoretical training and their use of intuition claims to confirm their own theoretical views.¹²

Philosophers are a rarefied and self-selected group, with a highly distinctive, largely shared educational history. It is not at all obvious that the intuitions they report will offer an accurate representation of the intuitions of non-philosophers. As Stephen Stich and Jonathan Weinberg put it, philosophers are not justified in assuming without any supporting evidence beyond their own intuitions and introspective hunches, 'that the intuitions of high socio-economic status males, like themselves, who have advanced degrees in philosophy, and whose cultural background is Western European can serve as a basis for generalizations about the intuitions of "the folk"' (2001, 642; see also Stich 1998). For similar reasons, it is also unclear to us how Kauppinen's suggested method of 'dialogue and reflection' avoids these worries, especially since he suggests that, 'in practice, assessing the truth of intuition claims can remain a relatively armchair business' (2007, 000). We take it that the degree to which philosophers' own intuitions and their reports of ordinary intuitions correlate with laypersons' actual intuitions is an empirical question and that it is amenable to controlled and systematic examination.

Indeed, the second assumption shared by experimental philosophers is that both ordinary language usage and the intuitions of laypersons—as well as the claims that philosophers make about them—can be approached through experimental methods (though, of course, not *only* through experimental methods). And they suggest that in cases where such methods have not been used by social scientists to examine these issues, philosophers with the requisite skills should do so themselves. So far, experimental philosophers have primarily used the methodology of presenting cases (sometimes derived from philosophical thought experiments) to non-philosophers and systematically collecting layperson's responses to questions about those cases. The statistically analyzed results of these studies have often been both surprising and perplexing. This partly explains one of the best features of the experimental philosophy movement—namely that even

philosophers who have competing explanatory models of the data frequently work together to develop the next round of experiments or to figure out the best way to interpret the results from all of the relevant studies about a given topic.

Despite the fact that all of the highly collaborative and interdisciplinary work that has been done by experimental philosophers so far has been driven by these overarching commitments, there are at least three distinct kinds of projects that have been undertaken. While some work in experimental philosophy seeks primarily to find out *what* the folk think in order to ascertain which philosophical theories best accord with and account for commonsense intuitions, other work primarily aims to explore *how* the folk think—what psychological mechanisms produce people’s intuitions or theories—and to determine the relevance of this information to philosophical disputes. Finally, a third project within experimental philosophy relies on data concerning cognitive diversity to argue that philosophers should not use intuitions as evidence in their theorizing. In the remainder of this section we will explain these three projects of experimental philosophy in more detail, after which we will turn our attention to Kauppinen’s criticisms.¹³

According to the first branch of experimental philosophy—which we are going to call ‘Experimental Analysis’ (EA)—the primary goal is to explore in a controlled and systematic manner what intuitions ordinary people tend to express and examine their relevance to philosophical debates.¹⁴ Hence, they aim to test philosophers’ claims that their positions align with common sense and to challenge those claims that are not supported by the evidence. On this view, philosophical theories that most closely accord with and account for ordinary beliefs and practices should enjoy ‘squatters’ rights’ until they are shown to be defective for other reasons. In this respect, EAs essentially agree with many traditional philosophers (including Kauppinen and those philosophers he discusses in Section 1 of his paper) about the *relevance* of folk intuitions; they simply disagree about the best *methods* for getting at these intuitions.¹⁵ Whereas the traditional philosopher is content to rely on armchair speculation, introspection, and informal dialogue, EAs utilize methods borrowed from experimental psychology to probe folk intuitions in a controlled and systematic way.

Of course, EAs realize that even if a philosophical theory is shown to be *more intuitive* than rival theories, it would not mean that the theory is *true*. Various responses are always open to philosophers whose theory has less support from such intuitions. For instance, they can offer *error theories* to explain away the intuitions (Greene 2002, forthcoming; Mackie 1977), or they can argue for *revision* of ordinary beliefs and practices in light of, for instance, our best current science (Vargas 2005). However, *until we first determine what the relevant intuitions actually are, any talk about error theories and revisionism is premature; indeed, experimental philosophy can play an important role in determining precisely when a theory is revisionist and when error theories are required.*

In addition to reporting the results of their studies, experimental philosophers also explore background issues such as the nature and sources of intuitions, the role that they should play in philosophy, how best to explore them, and what responses are available to theorists whose views do not settle with folk intuitions. Take, for example, our work with collaborators Jason Turner and Stephen Morris on folk intuitions about free will (Nahmias et al. 2005, 2006). Our initial motivation was to test the common claim that folk intuitions favor incompatibilism and that compatibilism is counter-intuitive.¹⁶ However, when our participants read about agents performing actions in a deterministic world, the majority judged that the agent does act of his own free will and is morally responsible for his

actions. We take these results to put pressure on the claim that incompatibilism is intuitive. But we also offer arguments for why data like ours are relevant to the free will debate more generally, and we discuss various incompatibilist responses to our methodology and our results. We point out, for instance, that incompatibilists may argue that it is the inferential *principles* leading to the incompatibilist conclusion that are intuitive, not the conclusion itself; if so, it would be helpful to determine if those principles are in fact intuitive to non-philosophers, since philosophers disagree about this question as well. Similarly, other work in experimental philosophy such as Joshua Knobe's research on the folk concept of intentional action (2003, 2004a, 2004b) has included and inspired numerous arguments about how and why the data are significant to debates about intentional action and how best to interpret the data in light of various philosophical theories.¹⁷

Of course, experimental philosophers need not limit themselves to exploring what folk intuitions are and why they matter to philosophy. They are also often interested in the *sources* of these intuitions. Questions that arise in this context include: What kind of cognitive mechanisms produce people's intuitions about particular cases and are these mechanisms reliable? How do we know which intuitions to trust, especially when we discover intuitions that conflict both within and between individuals? Might we better understand the persistence of some philosophical debates by understanding how different cognitive mechanisms generate conflicting intuitions? For another branch of experimental philosophy, these questions have taken center stage.

According to this second project—which we are going to call 'Experimental Descriptivism' (ED)—it is not only important to investigate what folk intuitions actually are, but it is also important to try to determine how these intuitions are generated.¹⁸ **EDs explore human psychology by testing how various manipulations to scenarios influence the intuitions people express. One goal of this project is to better understand the nature of the underlying psychological processes and cognitive mechanisms that produce our intuitions and explore the relevance of this research to philosophical questions.** But EDs are not interested in these cognitive processes and mechanisms for their own sake—rather, they rely on the gathering data concerning how the mind works to support (or attack) first-order theories in philosophy. Unlike EAs, who use the data concerning folk intuitions as foundations for conceptual analyses of 'free will,' 'intentional action,' and the like, EDs use the evidence to show that certain philosophical theories do not comport with what we are learning about how the mind works.

Consider, for instance, the research being done by Joshua Greene and his collaborators at the crossroads of neuropsychology and moral philosophy (Greene 2002, 2003; Greene and Haidt 2002; Greene et al. 2004).¹⁹ While most of the existing work in experimental philosophy has relied on the survey methodology, **his work demonstrates that there is nothing to prevent experimental philosophers (perhaps in collaboration with cognitive scientists) from utilizing more sophisticated ways of probing folk psychology.** Greene and his collaborators use neural imaging studies (fMRI) to examine what happens in people's brains as they consider famous thought experiments from moral philosophy. They have found that different scenarios not only tend to elicit different intuitions, but also activate different cognitive processes as well. Moreover, Greene has gone on to argue that data from his studies on moral cognition not only support sentimentalist over rationalist theories of moral psychology, but they also purportedly provide us with reasons to prefer consequentialist theories over their deontological counterparts.

Greene's research program is indicative of the way experimental philosophy, as an interdisciplinary and collaborative enterprise, can examine the neural and psychological

sources of people's intuitions about philosophical issues. There is no consensus about the relevance of such information to the philosophical debates about these issues, but there is no doubt that obtaining such information is driving important debates *about* the relevance of such information to philosophy. And as we will see below, information about the psychological mechanisms underlying people's intuitive judgments can also be useful in mapping out when people's judgments are influenced by seemingly irrelevant factors and in developing error theories for unreliable intuitive judgments.

The possibility that people's intuitions, and the underlying cognitive mechanisms that produce them, vary according to the kind of scenarios people consider complicates the picture for EAs and for philosophers who engage in the sort of conceptual analysis and theorizing that refers to ordinary conceptual usage and intuitions. When the data on intuitions collected thus far by experimental philosophers is coupled with the gathering evidence that some very basic moral intuitions vary systematically across different cultures and socio-economic classes (Haidt 2001, 2003; Haidt, Koller, and Dias 1993; Nisbett 2003), this puts pressure on the evidentiary status of intuitions with respect to philosophical theory building. **After all, if our intuitions depend in large part on philosophically irrelevant factors, and we select our philosophical theories at least in part based on how well they settle with our intuitions, we may end up unjustifiably privileging our own intuitions and customs over those of others.**

Highlighting these issues by gathering relevant data has been the primary concern of a third project within experimental philosophy. This more critical kind of experimental philosophy—which Joshua Alexander and Jonathan Weinberg have called **'Experimental Restrictionism' (ER)—can be viewed as an empirically informed battle against the use (and abuse) of intuitions in philosophy.**²⁰ The main goal of ER is to show that some of the methods and techniques that philosophers working in the analytic tradition have taken for granted are threatened by the gathering empirical evidence concerning both the diversity and the unreliability of folk intuitions. More specifically, ERs argue that *if* our intuitions about a particular topic vary cross-culturally or socio-economically and we don't have independent grounds for privileging our own intuitions to those of others, these particular intuitions will be insufficient for philosophical theory building.²¹

The worry expressed by ERs is that in light of the gathering evidence concerning the wide-scale diversity of intuitions, intuition-driven philosophy ends up being both epistemologically xenophobic and intuitionally question-begging. For instance, Weinberg, Stich, and Nichols have run studies that suggest Westerners apply the concept of knowledge differently than East Asians (e.g. East Asians are more willing to judge that people have knowledge in Gettier-style cases) (Nichols, Stich, and Weinberg 2003; Weinberg, Stich, and Nichols 2001). Similarly, Edouard Machery et al. (2004) found that East Asians are less likely than Westerners to use a Kripkean theory of reference. The results from these studies suggest that without a (culturally neutral) error theory to explain why one culture's intuitions are mistaken, these cross-cultural differences put pressure on both the evidentiary status of intuitions and on the ability of reflective equilibrium to help us separate the good intuitions from the bad ones.²²

The critical stance adopted by the ERs with respect to the proper role (or lack thereof) played by intuitions in certain areas of philosophy is noteworthy to the extent that it highlights just how different the various projects within experimental philosophy can be. While EAs are busy trying to discover what is intuitive in order to establish which philosophical theories have the most *prima facie* intuitive appeal, ERs are busy trying to establish that

the gathering data concerning the systematic diversity of intuitions gives us reason to disabuse ourselves of our intuition-mongering philosophical heritage. In this respect, experimental philosophy is a house divided—though fully exploring these divisions is a task for another day.²³ For now, we just wanted to highlight the central similarities and differences between the various camps within experimental philosophy so that we will be in a better position to examine the merit and proper scope of the criticisms that have recently been put forward by Kauppinen (2007) and others.

3. Responding to Criticisms: the Future of Experimental Philosophy

In this section we will argue that none of the criticisms that have been put forward so far succeed at showing that exploring ordinary intuitions and concepts using experimental methods is entirely irrelevant to philosophy. In our view, experimental philosophers have both the philosophical arguments and the methodological tools to allay many of the concerns and criticisms that have been raised. One of our strategies for responding to these criticisms is nicely captured by the transposed motto ‘united we (may) fall, but divided we stand.’ That is, some of the criticisms we will examine are only relevant to some of the projects we’ve outlined while other criticisms rely on presuppositions embraced by one or more of the projects within experimental philosophy. Another strategy we take for responding to criticisms is to turn them on their head and suggest that properly addressing them requires *more*, not less, experimental research. We will begin with what we take to be one of the most common objections to experimental philosophy.²⁴

Upon first hearing of experimental philosophy, many philosophers conclude from the start that empirical data concerning folk intuitions are irrelevant to philosophical debates because the folk intuitions *themselves* are irrelevant to such debates. After all, scientists and mathematicians tend not to worry about whether their theories settle with the intuitions of laypersons. Why should philosophers be any different? On this view, even if our own ‘expert’ intuitions correctly come into play when we’re doing philosophy, the untutored and uninformed intuitions of the ‘person on the street’ have no similar role to play.²⁵ One response to this line of reasoning is that the proponent of such a view owes us a non-question-begging account for why philosophers’ intuitions are more reliable than the pre-theoretical intuitions of laypersons. As we saw earlier, as the evidence concerning heuristic biases, cognitive diversity, situationism, and the like continues to mount with respect to children, laypersons, and experts, the philosopher’s faith in her own intuitions seems less warranted. Rather than assuming without evidence that introspection and reflective equilibrium give us access to the truth, philosophers have the burden of demonstrating that their own intuitions and the methods they use to sort them out are not subject to the same unwanted factors that they claim influence their folk counterparts. **Given the gathering data on the dark underbelly of human cognition—beset as it is by automatic processes, morally irrelevant situational cues, and cognitive biases and prejudices—it is no longer enough for philosophers to simply *assume* that their own expert intuitions are not susceptible to the factors that influence and sometimes distort folk intuitions. Ironically, to establish that pre-philosophical folk intuitions should *not* be trusted and that philosophically informed intuitions *should* be trusted would require more, not less, experimental research,** at least some of which should be focused specifically on intuitions about philosophical issues.

Another way of responding to the critic of experimental philosophy who rejects the importance of folk intuitions from the outset is to simply concede that not all folk intuitions are philosophically interesting. On this moderate view, the philosophical relevance of folk intuitions will vary from topic to topic. However, just because some intuitions may not be relevant to philosophy, it does not follow that we should so hastily banish *all* folk intuitions. Owing to the inherently abstract nature of some areas of philosophy—e.g. the philosophy of mathematics—we may not care at all about pre-theoretical judgments and intuitions with respect to those areas. But for many other philosophical issues, folk intuitions do—and arguably should—play an important role. This is especially true when the relevant concepts are intimately bound up with our everyday moral views and practices, such as free will and intentional action (compare Kauppinen 2007, 114). According to the experimental analysisists (EAs), while we need not be entirely beholden to folk intuitions concerning free will, moral responsibility, intentional action, and the like, these intuitions ought to at least serve as constraints on our philosophical investigations.

The stance adopted by EAs towards folk intuitions is one that is shared by some of the critics of experimental philosophy. Kauppinen, for instance, asks ‘why should anybody care about what philosophers do if they just argued about their own inventions?’ (2007, 96). On his view, because folk intuitions are used at least in part to adjudicate between competing analyses, philosophers ought to be interested in the intuitions of non-partisans—i.e. individuals who are not invested in the philosophical debate that is under investigation. This shows that both Kauppinen and at least some of the experimental philosophers he criticizes agree that conceptual analysis often does and should focus on folk concepts. They also seem to agree that these concepts can be mapped by examining the way people express certain intuitions to relevant cases (see Kauppinen 2007). The difference is that Kauppinen argues that experimental philosophers simply cannot get at the relevant ordinary intuitions and concepts using ‘detached, non-participatory social scientific research methods’ (2007, 97).

On this view, EAs are correct that philosophers ought to pay close attention to folk intuitions; they are simply incorrect in assuming that so-called ‘quasi-observational’ surveys give us access to these intuitions. According to Kauppinen, the methods used by experimental philosophers can only explore what he calls *surface intuitions*, whereas philosophers engaged in conceptual analysis ought to be interested in exploring only *robust intuitions*—i.e. intuitions competent language users would express under ideal conditions and without being influenced by irrelevant factors (2007, 101).²⁶ Kauppinen argues that it is highly doubtful—if not impossible—that experimental philosophers can shed any light on robust folk intuitions with the detached and non-participatory methods of social science. Instead, we are told that the philosophically interesting intuitions can only be uncovered using a participatory method he calls *dialogue and reflection* (2007, 98).

Kauppinen’s argument seems to be based on the assumption that experimental philosophers must limit themselves to the sort of ‘quasi-observation’ that he criticizes throughout the paper. But as we will show throughout this section, this is simply not the case. By focusing on the various experimental avenues that are open to experimental philosophers, we will show that Kauppinen’s conception of experimental philosophy is unjustifiably narrow. But first we will briefly examine his own preferred method for doing conceptual analysis—namely the Dialogue and Reflection Model (DRM). More specifically, we are going to show that DRM can be implemented in two distinct yet related ways—namely either informally or in controlled and systematic fashion.

According to the Informal Dialogue and Reflection Model (IDRM), it suffices for the purposes of doing conceptual analysis that we engage in discussion with our colleagues, students, friends, and family. Our stored memories of these conversations can then serve as the background against which each philosopher engages in dialogue and reflection—with each participant to the conversation bringing to the table her own unique history with the concept in question. When doing conceptual analysis along the lines of IDRM, we work on our analyses not only one concept at a time, but also one conversation at a time. But there are a number of problems with IDRM. For instance, each philosopher's intuitions will be shaped by where she grew up, her educational history, the variety of people she normally converses with, the myriad environmental circumstances within which these conversations took place, and the like. As a result, two different philosophers may understandably end up with two very different ideas with respect to what most people actually think about a particular topic or concept—despite the fact that each philosopher has engaged in the very kind of dialogue and reflection that Kauppinen suggests is supposed to help shed light on the truth. These differences will surely affect their conversations with interlocutors as well as affecting their beliefs concerning what is intuitive. Furthermore, it should be expected that philosophers' own intuitions as well as their opinions about ordinary intuitions and concepts have been influenced by the theoretical positions they have come to adopt during their extensive and particular philosophical training. Part of the problem may be that conflicting pre-philosophical intuitions drive philosophers to adopt competing theories, and philosophers with competing theories then highlight the intuitions that support their views.

Owing to these shortcomings of IDRM, proponents of Kauppinen's basic DRM may be inclined to adopt a different strategy for getting at the requisite data concerning folk intuitions and concepts. According to this second approach—which we are going to call Experimental Dialogue and Reflection Model (EDRM)—we should set up controlled and systematic experiments in order to find out what people's reflective intuitions, judgments, and beliefs about a given topic really are.²⁷ By having the conversations in a controlled environment, we could then code, compare, and analyze participants' answers in a rigorous and systematic manner that is less subject to the problems associated with IDRM. Moreover, we could collect the data markedly faster than we could by relying exclusively on informal conversations. With graduate students to assist in the collection and analysis of the data and using large numbers of participants in the studies, we could collect in a short amount of time information that it might take a single philosopher years to accumulate. As more studies are run and more data are collected, meta-analyses could eventually be run to see what patterns emerge both within and between large and diverse populations. By implementing the dialogue and reflection model in a controlled and systematic way, we could get at the data Kauppinen claims that we want in a manner that is less methodologically problematic than the informal method he advocates. And while implementing EDRM would admittedly require more time and resources in the beginning than implementing IDRM, once it is up and running it would enable us to access folk intuitions far more rapidly than informal conversation and dialogue alone.²⁸

By highlighting the possibility that an experimental approach could be used in implementing Kauppinen's own DRM, we learn three important lessons. First, we see that experimental philosophers have resources to get at what Kauppinen takes to be the philosophically interesting intuitions without having to simply retreat to the armchair in defeat. Indeed, they could get at these intuitions with a modified and less problematic

version of Kauppinen's own preferred method. Second, EDRM makes clear that any information that an individual philosopher can access via a series of informal conversations, experimental philosophers can access via controlled experiments that enable them to amass far more data over the long haul in far less time. Finally, it shows that experimental philosophers can, and in some cases should, be interested in more than the mere 'quasi-observational' method Kauppinen criticizes. And while they may have to borrow or develop more sophisticated methods in order to get at the data they really want (remember this is still a nascent movement), there is nothing preventing them from making these methodological modifications in order to advance the goals of their various projects (see below for more details).

In order to undermine the entire movement of experimental philosophy, Kauppinen would need to establish either (i) **that we simply cannot get at philosophically interesting intuitions regardless of what experimental designs we use, or (ii) even if we could get at the salient intuitions experimentally, doing so is not worth the effort.** As far as (ii) is concerned, we believe that probing the intuitions in a controlled and systematic way is worth the effort given the problems associated with relying exclusively on informal means of gathering information about conceptual usage and folk intuitions. At a minimum, before we can judge the merit of (ii), we would first need to test the reliability of the informal methods upon which philosophers have tended to rely. As far as (i) is concerned, we believe this view commits Kauppinen to an indefensibly strong modal claim. It is one thing to claim that the kinds of studies that have been run *so far* by experimental philosophers have not shed light on the salient folk intuitions—a claim we will reject below—it's quite another thing to claim that it is *impossible* that controlled and systematic experiments of any sort could do so.

Historically, claims that empirical methods could *never* achieve certain goals have withered under the progress of science and technology. As such, we think Kauppinen is incorrect to suggest that experimental philosophers have the burden of demonstrating that the methods of experimental psychology can give us access to philosophically interesting folk intuitions. Rather, the critic has the burden of showing that these intuitions lay forever beyond the reach of experimental philosophy. And one who wants to maintain the relevance of facts about folk intuitions walks a fine line in arguing that they *can* be properly discovered through non-experimental methods but they simply *cannot* be properly examined using any experimental methods.

As it stands, Kauppinen's arguments rest on the questionable assumption that the fate of experimental philosophy is tied to the currently used methods. Once we disabuse ourselves of that assumption, many of Kauppinen's objections to experimental philosophy lose much of their force. For instance, Kauppinen writes that experimental philosophy 'rests on mistaken assumptions about the relations between people's concepts and their linguistic behaviour' (2007, 95). And he suggests that practitioners of experimental philosophy have neglected the fact that philosophers who appeal to ordinary intuitions are being 'elliptical,' since they really mean to reference robust intuitions, not those offered by incompetent users or by competent users under problematic conditions.

As we saw earlier, Kauppinen believes that philosophers should not be interested in just any old folk intuitions, they should only be interested in robust intuitions—i.e. intuitions that satisfy three important conditions. First, in order for an intuition to be robust, its application must 'generally match the conceptual norms prevailing in the linguistic community' (2007, 103). Second, we are only interested in intuitions that competent language

users have under ideal conditions—i.e. ‘conditions in which there are no perturbing, warping or distorting factors or limits of information, access or ability’ (2007, 103). Finally, robust intuitions—unlike surface intuitions—are produced by semantic considerations rather than pragmatic considerations. According to Kauppinen, experimental philosophers are ill-equipped to adequately examine any of these three issues with the research methods available to them.

But if it can be shown that experimental philosophers can use controlled and systematic experiments to shed light on precisely these kinds of issues—and we will argue that they can—then we can turn Kauppinen’s criticisms on their head. Take, for example, his suggestion that with ‘surveys of amateurs it is practically impossible’ to ‘separate the contribution of semantic and pragmatic considerations to what people say (and what it is *proper* to say)’ (2007, 105). This is a surprising criticism for Kauppinen to make given the extended debate that has transpired among experimental philosophers concerning the very issue he raises. For instance, shortly after Knobe’s data first appeared, Fred Adams and Annie Steadman (2004a, 2004b) suggested that even though Knobe’s participants may have answered in a way that suggests on the surface that they judged that intending to *x* is unnecessary for intentionally *x*-ing, their answers were being driven by pragmatic considerations having to do with the relationship between intentional action and moral responsibility and not by semantic considerations having to do with the folk concept of intentional action.

Similarly, Alfred Mele (2001, 2003) developed an error theory to explain how and why moral considerations were affecting participants’ ascriptions of intentionality. On his view, because some participants *erroneously* assumed that if an agent is morally responsible for having *x*-ed, then the agent must have *x*-ed intentionally, they falsely ascribed intentionality to the agents in Knobe’s studies. And because it is ‘easy enough to show people that, upon consideration, they themselves would reject this assumption, their tendency to allow the moral features of actions to affect their ascriptions of intentionality is both mistaken and corrigible’ (Mele 2001, 41). Mele goes on to predict that by reminding participants of garden-variety cases of negligence and recklessness, they will then purportedly be able to see that the assumption that all blameworthy actions must have been performed intentionally is ‘false by their own lights rather than by the lights of an externally imposed theory’ (Mele 2001, 41). This should in turn lead them to give different answers with respect to intentionality.

Experimental philosophers have subsequently tested Mele’s prediction (Knobe 2003b; Nadelhoffer 2004a). In these studies, participants once again received the kind of vignettes Knobe used in his earlier side-effect studies—except this time participants were first primed with drunk-driving cases that were designed to make it clear that agents can be morally responsible for killing their victims even if they do not bring about their victims’ deaths intentionally. Nearly all the participants did indeed judge that the drunk driver was blameworthy even though they did not judge that he brought about the victim’s deaths intentionally. Yet, their responses to the follow-up vignettes were consistent with the responses given by participants in earlier studies who had not been similarly primed. It appears that being reminded that agents can be responsible for *x*-ing even if they did not *x* intentionally does not change participants’ intuitions concerning intentional action. The results of these studies suggest that the pragmatic consideration that Adams and Steadman (2004a, 2004b), Mele (2001, 2003), and Hugh McCann (2005) have all identified does not appear to be driving participants’ intuitions in the way they

had assumed. And while now is not the time to pick sides in this ongoing debate, **the fact that both experimental philosophers and psychologists are actively trying to address the very issue Kauppinen raises suffices to show that he is wrong in assuming that experimental philosophers either have not or cannot attempt to get at intuitions that are driven by semantic rather than merely pragmatic considerations.**

A similar strategy is open for responding to Kauppinen's claim that the *normative* components of robust intuitions—which require both linguistic competency and ideal conditions—simply cannot be examined using experimental methods (2007, 107). As we have already seen, Kauppinen believes that when philosophers appeal to ordinary intuitions they are being 'elliptical,' since they really mean to reference robust intuitions, not those offered by incompetent users or by competent users under problematic conditions. And since experimental philosophers are purportedly only able to get at surface intuitions, we are told that their work is of little or no philosophical importance, even if the data they collect may turn out to be of interest to psychologists (2007, 108). Here, we suspect that Kauppinen misunderstands both the methods most experimental philosophers have been employing and their goals in using those methods.

One way for experimental philosophers to respond to this concern would be to simply 'go conditional.' By this we mean that they could explore experimentally the answers to questions of this form: *If conditions are such and such, then people express such and such intuitions* (or tend to 'say that X is C,' to use Kauppinen's language). Then, competing philosophers can argue about whether *those* conditions are indicative of biasing or pragmatic factors or, conversely, are requisite triggering or background features for elucidating competence with the concepts. This is one way to understand what Experimental Descriptivists (EDs) are doing. Philosophers with competing theories can stipulate, based on their theory, the relevant normative conditions for competency, and then EDs can test whether these conditions are satisfied by particular subjects and, if so, test what those subjects say in response to philosophical probes designed to test their intuitions. Perhaps we misunderstand Kauppinen's claim that considerations of competence entail that intuition statements cannot be experimentally tested 'without taking a stance on the concepts involved' (2007, 97), but there seems nothing problematic about a researcher's accepting a particular theory (e.g. about what it takes to demonstrate competence with a particular concept) and then experimentally testing that theory. **Indeed, this seems to be analogous to the scientific practice of developing predictions from one's theory and then experimentally testing them.**

We also reject Kauppinen's suggestion that experimental philosophers aim to obtain (and are only *able* to obtain) data about surface intuitions rather than robust intuitions. We believe, rather, that the *goal* of most experimental philosophers is to obtain robust data about robust intuitions, and that any failures to achieve this goal are the result of methodological tools in need of fine-tuning, not mistaken assumptions about what sort of intuitions, if any, are relevant to philosophical theorizing. And we believe that the methodological tools are less imperfect than Kauppinen suggests, as we will now explain. But, first, we should again point out that the goal of many Experimental Restrictionists is to show through experimental methods that in some cases there are no robust intuitions with respect to the target concept.²⁹ On this view, the intuitions of philosophers, just like those of laypersons, may be subject to factors that are purportedly irrelevant, such as the vagaries of culture, class, even the order of presentation of the cases, and regarding philosophical debates where this is the case, the evidentiary value of intuitive appeals is negligible.

Kauppinen suggests that when traditional philosophers make claims about what is intuitive or 'what we would say about X,' 'charity requires us to conceive of their claims' as including certain background assumptions—namely these claims are about what *competent users* of the concepts would say if they considered them in *sufficiently ideal* conditions without being influenced by *non-semantic considerations* (2007, 101). Well, charity would also suggest that experimental philosophers are similarly aiming to obtain data about ordinary intuitions and conceptual usage under these conditions—though often they aim also to elucidate precisely *when* these conditions are and are not met. Kauppinen argues that the survey method ends up testing (E*), the intuitive judgments of non-specialists who '(1') *appear to understand the question . . .* (2') *however they consider it in whatever conditions they find themselves in* and (3') *whatever kind of considerations influence their response*' (2007, 105, italics in original). Our response is that experimental philosophers have used and continue to develop methods to determine (i) whether participants in their experiments properly understand the questions, (ii) whether they have considered them carefully, and (iii) whether or not they are influenced by clearly irrelevant factors.

To begin with, experimental philosophers use statistical methods. One of the primary benefits of statistical analysis is that it allows you to examine correlations between manipulated factors, even in the face of expected effects of the 'noise' of other factors. Using sufficiently large sample sizes, you can show that the probability is extremely low that the relevant results obtained because of the irrelevant factors.³⁰ **A statistically significant result does not indicate that any individual participant's response was not due to inattention or confusion or a mischievous desire to mess up the experiment. But for many types of studies, it does indicate that it is highly unlikely that most participants were inattentive or confused or mischievous—or if they were, they'd all have to 'mess up' in a relatively cohesive way.**

And experimental philosophers attempt to avoid participants' being affected by certain irrelevant factors (e.g. Kauppinen's non-ideal conditions). For instance, they instruct participants to read the scenarios and questions carefully and they provide them ample time to complete the surveys (usually they are present to observe that the participants at least appear to be taking it seriously).³¹ But experimental philosophers generally go on to *control* for whether participants are following these instructions by using 'manipulation checks' that test whether participants have understood the scenarios they have read. Participants who miss these questions, perhaps because they were not properly engaged or competent enough readers, are excluded from the analysis of the results.

Do these precautions ensure that only engaged and competent participants are included? No, some may guess correctly on manipulation checks, but again, statistical analyses ensure that only wide-scale confusion in a particular direction of response should produce any significant results obtained. Multiple questions using similar concepts can test (within-subjects) whether participants are being consistent in expressing their intuitions. Repeating surveys with different scenarios and across different populations can test for between-subject consistency. And randomizing the order of cases and questions can control for whether participants are influenced by the order of presentation.

In some cases, experimental philosophers require participants to reason counterfactually. For instance, in our free will studies we ask them to consider a scenario about another universe (described as deterministic) and answer questions about agents' freedom and responsibility in that universe. In these cases, we ask the participants to answer the experimental questions based on what the scenario says is true of the alternate universe, *not* what they believe is actually true of our universe (sometimes we ask whether

they think what it says is true of our world and then remind them, regardless of that, to answer based on the assumption it *is* true of the universe described in the scenario). It is difficult to ensure that participants successfully reason counterfactually, but there is a body of research in psychology that suggests people are generally quite adept at counterfactual reasoning, and the manipulation checks help test for whether they succeed.³² Again, barring widespread mistakes of a particular type, statistical analyses help to ensure that individual mistakes get ‘washed out.’

Finally, experimental philosophers often ask participants to explain *why* they responded to questions the way they did. Usually, the experimental goal is to explore the folk’s *intuitions* about whether, for instance, a certain concept applies in a certain scenario, not the folk’s *theories* about *why* the concept applies or doesn’t, so these open-ended explanations are only solicited *after* they have offered their intuitive judgments.³³ But these explanations then allow experimental philosophers to glean important information about whether the participants are reading carefully, understanding the scenarios and questions, reasoning counterfactually, and drawing appropriate inferences. Such explanations also offer useful information for fine-tuning future surveys.

We are belaboring these methodological points in order to emphasize that experimental philosophers tend to be acutely aware of the potential for irrelevant factors and biases influencing people’s intuitions and judgments. Given the interdisciplinary interests of most experimental philosophers, they tend to be familiar with the vast literature in experimental psychology that examines these factors and explores when and how they influence people. Indeed, for experimental philosophers engaged primarily in the ED (Descriptivist) and ER (Restrictionist) projects, the methods of psychology are employed precisely in order to better understand the folk ‘psychology of philosophy’ by examining which features of cases influence people’s intuitions and which cognitive mechanisms might thereby get implemented along the way.

These philosophers (and of course, the psychologists who engage in similar projects) tend to use the classic experimental paradigm: expose two randomly selected groups to scenarios that differ only in one variable and ask them the same questions about the scenario; if the two groups’ responses differ in a statistically significant way, we can conclude that it is highly likely that *that* variable accounts for the difference in the responses. That is, something about our psychology is influenced by that variable. For instance, if Knobe is right that the moral dimension of the side-effect is the only relevant difference between the scenarios he presents, then it is clear that the way the folk understand intentional action is influenced by moral considerations (2003a, 2003b). Alternatively, experimentalists may expose two groups that differ in only one crucial respect (e.g. cultural background or education level) to the *same* scenarios; if the groups’ responses differ in a statistically significant way, we can conclude that the crucial difference between the groups is driving the different responses—i.e. something about our psychology regarding these cases is affected by our culture, education level, etc.

Like Kauppinen, those engaged in the Experimental Analysis (EA) project are less interested in the *differences* between people’s responses than the similarities that indicate some shared intuition or concept. That is, their goal is to empirically confirm or disconfirm the claims of philosophers who appeal to intuitions as some kind of support for their theory or conceptual analysis. So, rather than using controlled and systematic studies primarily to discover psychological factors that drive *differing* intuitions, EAs use these studies to obtain systematic results about how the majority of people respond to cases and to control for

problematic influences (non-ideal conditions) on these responses. Nonetheless, in almost all studies, even if there is a statistically significant majority responding in one way, there is still a non-trivial *minority* responding in other ways. What sense do we make of the competing intuitional judgments one finds both *within* groups and *between* groups?

This is one of the questions used to motivate a criticism of experimental philosophy that Kauppinen calls the ‘argument from disagreement.’³⁴ On his view, the experimental philosopher faces a dilemma:

Either a test subject’s response to a survey question reveals whether the case falls under her concept, or it doesn’t. If it doesn’t, the response is obviously uninformative and running the survey for this purpose is pointless. But what if it does? Then those who answer in the negative will not and cannot *disagree* with those who answer in the affirmative. (2007, 108)

Kauppinen goes on to suggest that in order to be able to ask intelligibly whether two individuals A and B agree (or disagree) concerning whether a particular case falls under a concept *x*, we must first presuppose that these individuals have a shared concept of *x*. But if all of the participants within a group do have a shared concept, then why don’t they answer the same way? If they do not have a shared concept, on the other hand, then we get the purportedly counterintuitive result that participants who give conflicting answers to a survey question do not actually disagree at all. Given the non-unanimous results of the studies that have been run so far by experimental philosophers—and the potential problems that this thereby creates—we agree with Kauppinen that experimental philosophers (especially EAs) need to explain why they tend to privilege the answers given by the majority of participants.

EAs sometimes focus on the answers given by the majority of participants because they are testing claims made by philosophers about which of two philosophical theories has *wider* intuitive appeal. Our own work on the free will debate serves as a good example of this approach (Nahmias et al. 2006). The assumption that we set out to test was the well-worn claim among philosophers that incompatibilism is the commonsensical view. Contrary to what Kauppinen suggests, we do not believe that these appeals to the pre-theoretical intuitions of laypersons are *always* shorthand for ‘what perfectly competent users would believe or intuit under ideal conditions.’ Consider, for instance, the following claim by Galen Strawson:

Almost all human beings believe that they are free to choose what to do in such a way that they can be truly, genuinely responsible for their actions in the strongest possible sense—responsible period, responsible without any qualification, responsible *sans phrase*, responsible *tout court*, absolutely, radically, buck-stoppingly responsible; ultimately responsible, in a word—and so ultimately *morally* responsible when moral matters are at issue. Free will is the thing you have to have if you’re going to be responsible in this all-or-nothing way. That’s what I mean by free will. That’s what I think we haven’t got and can’t have.³⁵

Here we find Strawson explicitly claiming that what he means by free will and moral responsibility is what ‘almost all human beings’ mean by it. It is because of claims like this that EAs sometimes focus on how the majority of participants respond to their surveys, and they need not assume that people share the same intuitions about certain concepts.

Experimental philosophers may also analyze *why* different participants offer different responses to philosophical surveys.

For instance, Nichols and Ulatowski (forthcoming) have recently performed some ingenious experiments on Knobe cases, the results of which suggest that there may be two concepts of intentional action, one of which focuses on the agent's intention/motive and the other on the agent's foreknowledge of the action's effects. By asking participants to explain their answers and then coding these answers to see what they have in common, Nichols and Ulatowski were better able to investigate *when* participants were employing one of these concepts rather than the other. For present purposes, their research serves as an important reminder that experimental philosophers need not focus only on how the majority of participants respond to thought experiments, especially when paying attention to the minority responses is fruitful as well.³⁶

As we learn more about what drives people's intuitions, we will also be in a better position to determine whether members in the minority (or even the majority) group are making particular kinds of mistakes. This would in turn allow us to develop error theories for why some people (or some groups) express certain intuitions that they would reject upon further reflection. But, once again, adequately exploring these issues will require *more rather than less experimental research*. Furthermore, we believe that in carrying out this research experimental philosophers can and should spend more time thinking about how to design studies so that they are better able to get at precisely the kinds of robust intuitions that Kauppinen mistakenly suggests are forever beyond their experimental grasp.

In our view, if the philosophically interesting intuitions can be accessed at all, we see no reason why they cannot be accessed using controlled and systematic methods that are likely to be less problematic than informal alternatives. And while getting at these intuitions may require more 'dialogue and reflection' on the part of participants than previous studies have used, we are unconvinced by Kauppinen's claim that experimental methods can have no possible role to play. We've already explained how experimental philosophers vary examples and ask for explanations (the first step in a dialogue). Now we will examine some other methods they have used or could use to address some of Kauppinen's concerns.

Several experimental philosophers have tried to see how people respond in the face of apparently conflicting intuitions, engaging them in a process in the vicinity of reflective equilibrium. For instance, Nichols and Knobe (forthcoming) present a universe (Universe A) they describe as deterministic and find that most participants respond that agents *cannot* be 'fully morally responsible' in that universe.³⁷ But when presented with a particular agent doing a dastardly deed in Universe A, most participants respond that he *can* be fully morally responsible. Nichols and Knobe then followed up with some of those participants who expressed both of these seemingly contradictory intuitions and asked them how they would resolve the conflict. That is, they entered into a (brief) dialogue with their participants to tease out some implications and see how they would respond.

The resources available for online surveys can be used in similar ways. In a large online study on intuitions about free will and responsibility, Nahmias, Coates, and Kvaran (forthcoming) had over 600 participants read a scenario and answer a dozen questions, which allows extensive statistical analyses of correlations between participants' responses to questions about free will, deliberation, moral responsibility, praise and blame, and retributive punishment, as well as demographic factors such as gender and degree of religiosity. We can examine apparent inconsistencies within individual participants' responses and connections between the ways people employ different concepts.³⁸ And of course, we

are excluding those participants who miss either of two manipulation checks, one administered before and one after the experimental questions. We will also examine whether the responses of those who have taken a college philosophy course differ from those who have not. Eventually, we may run studies to compare participants' responses before and after they have been taught the free will debate by philosophy teachers who hold opposing views to examine potential effects of theoretical biasing on pre-philosophical intuitions.

In another online survey, Nahmias and Sommers will ask participants for their level of agreement or disagreement first with the premises of something like Galen Strawson's 'Basic Argument' for skepticism about free will and moral responsibility (1994) and then with the conclusion of that argument (i.e. that it is impossible to be truly morally responsible for one's actions). The online program allows different participants to take different 'paths,' depending on their answers, so that individual participants will be offered various ways to resolve and try to explain any conflicts between their responses—roughly, to seek reflective equilibrium by being confronted, for instance, with the inconsistency between an initial acceptance of all of the premises and an initial rejection of the conclusion that validly follows from them.

This method of trying to track the thought processes of laypersons as they are introduced to philosophical arguments or thought experiments has the potential to test Experimental Analysts' hypotheses about what is most intuitive to the folk, Experimental Descriptivists' hypotheses about the psychological mechanisms underpinning the folk's intuitions, and Experimental Restrictionists' hypotheses about the instability of folk intuitions. If this method starts to look more like the Socratic dialogue Kauppinen advocates, that's fine, but it is decidedly *not* 'to do philosophy pretty much as it has always been done' (2007, 110).

Finally, the future of experimental philosophy, especially the ED branch, is likely to become increasingly interdisciplinary, involving collaboration between philosophers who understand the complexities of the philosophical debates and scientists who understand the complexities of the methodology required to shed light on those debates. Such collaboration might include behavioral experiments—to examine how people's beliefs about philosophical issues (e.g. free will or morality) influence their behavior.³⁹ It might include further fMRI studies—to examine the neural activity correlated with certain philosophical intuitions or moral beliefs. It might even include phenomenological investigations—to examine people's first-person reports about experiences relevant to philosophical questions, such as free will and deliberation (see Nahmias et al. 2004). When coupled with the other examples we have already discussed in this section, we believe these possibilities demonstrate that the future of experimental philosophy looks bright, so long as we don't assume from the start that the movement should fall before these and other possibilities are explored.

4. Conclusion

In this paper we have tried to shed light on the wide variety of projects within experimental philosophy while defending the movement from its critics. By reminding ourselves of the multitude of research methods open to experimental philosophers, we can better resist the limited vision of the experimental possibilities exhibited by critics such as Kauppinen. On our view, if we pay close attention to what experimental philosophers have actually been doing so far, while keeping our minds open with respect to how they might adopt new strategies in the future, we find that many of the very issues Kauppinen

raises can adequately be addressed via controlled and systematic experiments. Indeed, experimental philosophers can examine when intuitions are being expressed in less than ideal conditions, when they are being driven by semantic rather than pragmatic factors, and when they are being expressed by competent users.

While we readily acknowledge that Kauppinen has raised some important concerns, none of them requires experimental philosophers to rest content with the limitations of armchair reflection and informal dialogue. Moreover, we believe that determining the merit of several of his objections—as well as the merit of his own dialogue and reflection model—actually requires *more* rather than *less* experimentation. We are also aware that evaluating Kauppinen's view requires more rather than less philosophical discussion about issues we have not delved into here, such as the nature of intuitions and the proper analysis of concepts. Indeed, we appreciate that Kauppinen has highlighted a number of issues that experimental philosophers would do well to address more carefully. But given the surprising and informative results that experimental philosophers have produced so far and the interesting dialogue their work has already inspired, we are convinced that it is a mistake to either call for or predict the downfall of experimental philosophy.

ACKNOWLEDGEMENTS

For helpful comments, we would like to thank Joshua Alexander, Adam Feltz, Antti Kauppinen, Joshua Knobe, Trevor Kvaran, Shaun Nichols, and Jonathan Weinberg.

NOTES

1. Authorship is equal.
2. *Action Theory*: Adams and Steadman (2004a, 2004b), Knobe (2003a, 2003b, 2004b, 2005a, 2005b), Knobe and Burra (2006a, 2006b), Knobe and Mendlow (2004), Leslie, Knobe, and Cohen (2006), Machery (forthcoming), Malle (2001, 2006), Malle and Knobe (1997), McCann (2005), Meeks (2004), Nadelhoffer (2004a, 2004b, 2005, 2006a), Nichols and Ulatowski (forthcoming), Phelan and Sarkissian (forthcoming), and Young et al. (2006). *Epistemology*: Alexander and Weinberg (2006), Baron (1998), Bishop and Trout (2005a, 2005b), Feltz (n.d.), Nichols, Stich, and Weinberg (2003), Stich and Nisbett (1980), Stich and Weinberg (2001), Swain, Alexander, and Weinberg (forthcoming), and Weinberg, Nichols, and Stich (2001). *Ethics*: Cushman, Young, and Hauser (2006), Doris (2002), Greene (2002, 2003, forthcoming), Greene and Haidt (2002), Greene et al. (2004), Haidt (2001, 2003), Haidt, Bjorklund, and Murphy (2000), Hauser (2006), Hauser, Young, and Cushman (forthcoming), Knobe (2005a), Knobe and Leiter (forthcoming), Knobe and Roedder (forthcoming), Nichols (2002, 2004b, 2004c), Nichols and Mallon (2006), Pizarro and Bloom (2003), Pizarro, Uhlmann, and Bloom (2003), and Pizarro, Uhlmann, and Salovey (2003). *Free Will & Moral Responsibility*: Knobe and Doris (forthcoming), Nadelhoffer and Feltz (forthcoming), Nahmias (2006), Nahmias et al. (2005, 2006, forthcoming), Nichols (2004c, 2006a, 2006b), Nichols and Knobe (forthcoming), Turner and Nahmias (2006), Woolfolk and Doris (2002), and Woolfolk, Doris, and Darley (forthcoming). *Philosophy of Language*: Machery et al. (2004), and Mallon et al. (forthcoming). *Philosophy of Law*: Carlsmith, Darley, and Robinson (2002), Greene and Cohen (2004), Knobe (2005b), Malle and Nelson (2003), Nadelhoffer (2006a), and Sardjeveladzé and Machery (n.d.). *Philosophy of Mind*: Gray, Gray, and Wegner (2007), and Knobe and Prinz (forthcoming).

Philosophy of Science: Stotz (forthcoming), Stotz and Griffiths (2004), and Stotz, Griffiths, and Knight (2004).

3. We will speak primarily of ordinary (or 'folk') intuitions, which we take roughly to be **dispositions to make certain non-deductive, spontaneous judgments, for instance, about whether or not a particular concept applies in a particular situation** (see Goldman and Pust 1998; Jackson 1998). Some experimental philosophy, however, may be better understood as investigating folk *theories* or folk *concepts*, and the complicated connections between intuitions, theories, and concepts adds a layer of complexity both to what experimental philosophers take themselves to be doing and what their critics take them to be doing.
4. In fact, we think the criticisms of experimental philosophy that have been put forward so far (e.g. Kauppinen 2007; Levy n.d.; Sosa 2005, 2007) are philosophically interesting enough that experimental philosophers deserve some credit for having inspired them! We also think it is essential for experimental philosophers themselves to examine these metaphilosophical questions as they present their research so that they are as clear as possible about how they think their own research bears on philosophical debates.
5. For the purposes of this paper, we need not try to resolve the tension that exists between the various projects within experimental philosophy. Nor do we need to take sides. Our present goal is just to identify and briefly explore three fairly distinct projects within experimental philosophy.
6. The boundary is blurry (i) because most experimental philosophers are happy to obtain the data they take to be relevant to the philosophical issues they are exploring by drawing on already existing scientific research, and (ii) because some scientists are interested in the same kind of intuitions that interest philosophers and they, too, discuss the philosophical implications of their research (e.g. Baron 1998; Damasio 1994; Greene 2002, 2003, forthcoming; Haidt 2001, 2003; Hauser 2006; Hauser et al. forthcoming; Malle 2001, 2006; Nisbett 2003; Nisbett and Ross 1980; Wegner 2002). Because 'experimental philosophy' is perhaps best viewed as a family resemblance term, trying to explicate the movement in terms of necessary and jointly sufficient conditions would be wrong-headed even if well-intentioned. But the following two questions are nevertheless helpful. **First, do you run controlled and systematic studies and use the resultant data to shed light on philosophical problems? Second, do you sometimes address the tension that exists between what philosophers say about intuition and human cognition, on the one hand, and what researchers are discovering about these things, on the other hand?** Experimental philosophy can be seen as a movement at the cross-roads of psychology and philosophy that is being driven primarily by people who answer these two questions in the affirmative.
7. See, e.g. Jackson (1998), Jackson and Pettit (1995), Lewis (1972), and Smith (1994).
8. This shared distrust in armchair reflection and informal dialogue as generally reliable guides to folk intuitions is roughly what Kauppinen describes as 'EXPERIMENTALISM –' (2007, 000).
9. See, e.g. Mynatt et al. (1977), Nickerson (1998), and Wason (1960).
10. See, e.g. Ditto and Lopez (1992), and Edwards and Smith (1996).
11. See, e.g. Fields and Schuman (1976), and Ross, Greene, and House (1977).
12. We are not claiming that philosophers *are* in fact more susceptible to these well-known biases—that, of course, is an empirical question that calls for controlled and systematic investigation. Rather, we are merely suggesting that given the central role that many

philosophers give to folk intuitions, there is reason to worry about what Alvin Goldman and Joel Pust call ‘theory contamination’ (1998, 183). Minimally, we believe that the gathering data on cognitive biases place the burden clearly at the feet of philosophers who claim their own intuitions are a reliable guide to what people generally find intuitive.

13. The three projects we describe are not meant to label individual *philosophers* but rather individual *research projects* or *articles*. Some philosophers, such as Shaun Nichols, have worked on all three of the projects we will discuss.
14. What we are calling EA is roughly equivalent to what Kauppinen calls ‘optimistic experimentalism’ (2007, 99), and what Alexander and Weinberg call the ‘proper foundations view’ (2006). Some examples of EA include Knobe (2003a, 2003b, 2004a, 2004b), Nadelhoffer (2005, 2006a, 2006b), Nahmias et al. (2005, 2006, forthcoming), Nichols (2004a, 2004b), and Nichols and Ulatowski (forthcoming).
15. As we will see, not only is there no consensus among philosophers in general concerning the proper role (if any) that folk intuitions have to play in philosophy, but experimental philosophers are themselves divided on this issue.
16. Robert Kane exemplifies this view with his claim that ‘most ordinary persons start out as natural incompatibilists. They believe there is some kind of conflict between freedom and determinism . . . Ordinary persons have to be talked out of this natural incompatibilism by the clever arguments of philosophers’ (1999, 217).
17. See, e.g. Adams and Steadman (2004a, 2004b), Cushman, Young, and Hauser (2006), Machery (forthcoming), Malle (2006), McCann (2005), Meeks (2004), Nadelhoffer (2004b, 2005), Nichols and Ulatowski (forthcoming), and Phelan and Sarkissian (forthcoming).
18. Some examples of ED work in experimental philosophy include Greene (2002, 2003, forthcoming), Greene and Haidt (2002), Greene et al. (2004), Knobe and Doris (forthcoming), Nadelhoffer (2006a), and Nichols and Knobe (forthcoming). Knobe’s (2007) response to Kauppinen (2007) offers a more extended defense of the philosophical relevance of at least this branch of experimental philosophy, arguing that it advances our understanding of human nature in the ways traditional philosophers have attempted to do since the Greeks.
19. Greene is an interesting case as he is appointed in a psychology department despite having done his graduate work in philosophy. He is also one of only a handful of psychologists who self-identifies with the experimental philosophy movement.
20. What we are calling ER maps roughly onto what Kauppinen calls ‘pessimistic experimentalists’ (2007, 98). For some examples of ER projects, see Alexander and Weinberg (2006), Machery et al. (2004), Mallon et al. (forthcoming), Nichols, Stich, and Weinberg (2003), and Weinberg, Nichols, and Stich (2001).
21. ER is heavily rooted in the pioneering work begun by Stich in the early 1980s. Indeed, not only did Stich pave the way for ER, he also participated in some of the first studies in experimental philosophy. See, e.g., Nichols, Stich, and Weinberg (2003), Stich (1983, 1984, 1985, 1990, 1994, 1996, 1998), Stich and Nisbett (1980), Stich and Weinberg (2001), and Weinberg, Nichols, and Stich (2001).
22. Another project within the ER camp is to show that seemingly irrelevant features significantly influence the intuitions people express. For instance, the order in which cases are presented seems to influence people’s epistemic intuitions (see Swain, Alexander, and Weinberg forthcoming).
23. See Nadelhoffer and Nahmias (n.d.).
24. For the purposes of this paper, we neither pursue nor examine various other responses experimental philosophers have already offered to critics. Knobe (forthcoming), for

instance, suggests that even *if* Kauppinen were correct that experimental philosophers don't (and can't) get at the intuitions philosophers need for doing conceptual analysis, experimental philosophy is nevertheless useful for philosophers who are not interested in doing conceptual analysis but who are interested instead in the naturalistic project of understanding what the mind is and how it works. Another important issue that we will leave unaddressed is that many of the criticisms that have been put forward so far by Kauppinen, Sosa, and others rely too heavily on assumptions that are driven by unsettled issues in linguistics and the philosophy of language. As Machery (forthcoming) has correctly pointed out, until we have a better empirical understanding of the nature of concepts and how they are used, we simply cannot resolve several of the ongoing debates both about and within experimental philosophy concerning intuitions and conceptual analysis. We think Machery may be largely correct about this. But even so, as he goes on to suggest, experimental philosophers can nevertheless continue to fruitfully investigate the nature and limitations of human cognition. We believe that when the arguments put forward in the present paper are viewed in conjunction with the arguments developed by Knobe, Machery, Alexander and Weinberg, and others, we find a multi-pronged defense of experimental philosophy that successfully staves off all would-be assassins.

25. See Bealer (1996, 1998, 2000), Levy (n.d.), and Williamson (2004).
26. A similar point has been developed by Ernest Sosa—see, e.g. Sosa (2005, 2007). A response to Sosa can be found in Alexander and Weinberg (2006).
27. Some graduate students at Florida State University recently ran a pilot study that involved presenting participants with various cases about intentional action and allowing them to discuss and debate the cases among themselves. At the end of the study they took further surveys to examine how the students' views changed (or did not change). And while the results were inconclusive—owing primarily to some problems with the design of the studies—their strategy is certainly one that could prove useful in the future.
28. This is clearly not to suggest that philosophers should quit having informal conversations with their children, families, friends, students, and colleague about philosophical issues! Nor is to suggest that all of the information that we glean from these conversations is useless or uninteresting. We are simply cautioning against the dogmatic assumption that there is data that we can best (or only) obtain in informal one-on-one conversations.
29. Even in cases where people express consistent intuitions across cases, ERs argue that if people in different cultures or classes express consistently different intuitions, then picking out which people get to count as competent will often beg the question at issue. Hence, robust intuitions *within* a group may not satisfy Kauppinen's criterion of competence for robust intuitions.
30. A statistically significant result is one for which the probability that the result occurred due to chance factors was less than 5 percent—a *p* value of <0.05 indicates this probability.
31. Contrast this with the way most in-class 'surveys' of students' intuitions are conducted: the teacher describes the thought experiment or scenario verbally, usually once quickly, and then asks students to raise their hands on the spot to indicate their intuitions about it. See Nadelhoffer and Nahmias (forthcoming) for an explanation of the usefulness of the controlled survey method for *teaching* philosophy in the classroom.
32. See, e.g. German and Nichols (2003), and Harris, German, and Mills (1996).
33. Some argue that *all* intuitions are suffused with background theoretical beliefs. As we mentioned earlier, the connections between intuitions and theories (and concepts) are

complex, and we suspect experimental philosophy can help provoke philosophical debates about these connections, and presumably it can gather empirical data relevant to such debates.

34. A similar objection has been put forward by Ernest Sosa (see, e.g. 2005, 2007).
35. In a 2003 interview with Tamler Sommers for *The Believer* (http://www.naturalism.org/strawson_interview.htm).
36. Another explanation for divergent results is that some concepts may not have clear necessary and sufficient conditions (or at least, not the folk's concepts). In such cases, people may have shared intuitions about the paradigmatic applications of the concepts, but differing intuitions about the boundaries, as demonstrated by their different judgments about certain, often non-paradigmatic, cases.
37. We have worries about the way Nichols and Knobe describe determinism that make us skeptical that they have shown most ordinary people have incompatibilist intuitions (see Nahmias 2006; Turner and Nahmias 2006), but we think the differences they find across scenarios are fascinating and important.
38. For instance, an initial look at the data suggests that many people answer differently to questions about (i) free will compared to moral responsibility and (ii) praise compared to blame, while (iii) people answer almost identically to questions about free will compared to questions about whether decisions are 'up to' the agent.
39. See Vohs and Schooler (n.d.) who show that inducing people to doubt the existence of free will increases their willingness to cheat on a test.

REFERENCES

- ADAMS, F., and STEADMAN, A. 2004a. Intentional action in ordinary language: Core concept or pragmatic understanding. *Analysis* 64: 173–81.
- . 2004b. Intentional actions and moral considerations: Still pragmatic. *Analysis* 64: 264–67.
- ALEXANDER, J., and WEINBERG, J. 2006. Analytic epistemology and experimental philosophy. *Philosophy Compass* 2 (1): 56–80.
- BARON, J. 1998. *Judgment misguided: Intuition and error in public decision making*. Oxford: Oxford University Press.
- BEALER, G. 1996. A priori knowledge and the scope of philosophy. *Philosophical Studies* 81: 121–42.
- . 1998. Intuition and the autonomy of philosophy. In *Rethinking intuition: The psychology of intuition and its role in philosophical inquiry*, edited by M. DePaul and W. Ramsey. Lanham, Md.: Rowman and Littlefield.
- . 2000. A theory of the a priori. *Pacific Philosophical Quarterly* 81 (1): 1–30.
- BISHOP, M., and J. D. TROUT. 2005a. The pathologies of standard analytic epistemology. *Nous* 39 (4): 696–714.
- . 2005b. *Epistemology and the psychology of human judgment*. Oxford: Oxford University Press.
- CARLSMITH, K., J. DARLEY, and R. ROBINSON. 2002. Why do we punish? Deterrence and just deserts as motives for punishment. *Journal of Personality and Social Psychology* 83 (2): 284–99.
- CUSHMAN, F. A., L. YOUNG, and M. D. HAUSER. 2006. The role of reasoning and intuition in moral judgments: Testing three principles of harm. *Psychological Science* 17 (12): 1082–89.
- DAMASIO, A. 1994. *Descartes' error: Emotion, reason, and the human brain*. London: Penguin Books.

- DITTO, P. H., and D. F. LOPEZ. 1992. Motivated skepticism: Use of differential decision criteria for preferred and non-preferred conclusions. *Journal of Personality and Social Psychology* 63: 568–84.
- DORIS, J. 2002. *Lack of character: Personality and moral behavior*. Cambridge: Cambridge University Press.
- EDWARDS K., and E. E. SMITH. 1996. A disconfirmation bias in the evaluation of arguments. *Journal of Personality and Social Psychology* 71: 5–24.
- FELTZ, A. (n.d.) Problems with Bealer's appeal to intuition in epistemology. Unpublished manuscript.
- FIELDS, J., and H. SCHUMAN. 1976. Public beliefs about the beliefs of the public. *Public Opinion Quarterly* 40: 427–48.
- GERMAN, T., and S. NICHOLS. 2003. Children's counterfactual inferences about long and short causal chains. *Developmental Science* 6: 514–23.
- GOLDMAN, A., and J. PUST. 1998. Philosophical theory and intuitional evidence. In *Rethinking intuition: The psychology of intuition and its role in philosophical inquiry*, edited by M. DePaul and W. Ramsey. Lanham, Md.: Rowman and Littlefield.
- GRAY, H., K. GRAY, and D. WEGNER. 2007. Dimensions of mind perception. *Science* 315: 219.
- GREENE, J. 2002. The terrible, horrible, no good, very bad truth about morality and what to do about it. Ph.D. diss., Princeton University.
- . 2003. From neural 'is' to moral 'ought': What are the moral implications of neuroscientific moral psychology. *Neuroscience* 4: 847–50.
- . Forthcoming. The secret joke of Kant's soul. In *Moral Psychology, Vol. 3: The neuroscience of morality: Emotion, disease, and development*, edited by W. Sinnott-Armstrong. Cambridge, MA: MIT Press.
- GREENE, J. D., and J. D. COHEN. 2004. For the law, neuroscience changes nothing and everything. *Philosophical Transactions of the Royal Society of London B* 359: 1775–1785.
- GREENE, J., and HAIDT, J. 2002. How (and where) does moral judgment work? *Trends in Cognitive Science* 6 (12): 517–23.
- GREENE, J., L. NYSTROMM, A. ENGELL, J. DARLEY, and J. COHEN. 2004. The neural bases of cognitive conflict and control in moral judgment. *Neuron* 44: 389–400.
- HAIDT, J. 2001. The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review* 108: 814–34.
- . 2003. The emotional dog does learn new tricks: A reply to Pizarro and Bloom. *Psychological Review* 110 (1): 197–98.
- HAIDT, J., F. BJORKLUND, and S. MURPHY. 2000. Moral dumbfounding: When intuition finds no reason. Unpublished manuscript.
- HAIDT, J., S. KOLLER, and M. DIAS. 1993. Affect, culture, and morality, or is it wrong to eat your dog? *Journal of Personality and Social Psychology* 65: 613–28.
- HARRIS, P. L., T. P. GERMAN, and P. E. MILLS. 1996. Children's use of counterfactual thinking in causal reasoning. *Cognition* 61: 233–59.
- HAUSER, M. D. 2006. *Moral minds*. New York: Harper Collins.
- HAUSER, M. D., F. A. CUSHMAN, L. YOUNG, R. KANG-XING JIN, and J. MIKHAIL. Forthcoming. A dissociation between moral judgments and justifications. *Mind & Language*.
- HAUSER, M. D., L. YOUNG, and F. A. CUSHMAN. Forthcoming. Reviving Rawls' linguistic analogy. In *Moral psychology and biology*, edited by W. Sinnott-Armstrong. New York: Oxford University Press.

- JACKSON, F. 1998. *From metaphysics to ethics: A defense of conceptual analysis*. Oxford: Oxford University Press.
- JACKSON, F., and P. PETTIT. 1995. Moral functionalism and moral motivation. *Philosophical Quarterly* 45: 20–40.
- KANE, R. 1999. Responsibility, luck, and chance: Reflections on free will and indeterminism. *Journal of Philosophy* 96: 217–40.
- KAUPPINEN, A. 2007. The rise and fall of experimental philosophy. *Philosophical Explorations* 10 (2): 95–118 (this issue).
- KNOBE, J. 2003a. Intentional action and side-effects in ordinary language. *Analysis* 63: 190–93.
- . 2003b. Intentional action in folk psychology: An experimental investigation. *Philosophical Psychology* 16 (2): 309–23.
- . 2004a. Folk psychology and folk morality: Response to critics. *Journal of Theoretical and Philosophical Psychology* 24 (2): 270–79.
- . 2004b. Intention, intentional action and moral considerations. *Analysis* 64: 181–87.
- . 2005a. Theory of mind and moral cognition: Exploring the connections. *Trends in Cognitive Sciences* 9: 357–59.
- . 2005b. Cognitive processes shaped by the impulse to blame. *Brooklyn Law Review* 71: 929–37.
- . 2007. Experimental philosophy and philosophical significance, *Philosophical Explorations* 10 (2): 95–118 (this issue).
- . Forthcoming. The concept of intentional action: A case study in the uses of folk psychology. *Philosophical Studies*.
- KNOBE, J., and A. BURRA. 2006a. The folk concept of intention and intentional action: A cross-cultural study. *Journal of Cognition and Culture* 6 (1/2): 113–32.
- . 2006b. Experimental philosophy and folk concepts: Methodological considerations. *Journal of Cognition and Culture* 6 (1/2): 331–42.
- KNOBE, J., and J. DORIS. Forthcoming. Strawsonian variations: Folk morality and the search for a unified theory. In *The handbook of moral psychology*, edited by J. Doris et al. Oxford: Oxford University Press.
- KNOBE, J., and B. LEITER. Forthcoming. The case for Nietzschean moral psychology. In *Nietzsche and morality*, edited by Brian Leiter and Neil Sinhababu. Oxford: Oxford University Press.
- KNOBE, J., and G. MENDLOW. 2004. The good, the bad, and the blameworthy: Understanding the role of evaluative reasoning in folk psychology. *Journal of Theoretical and Philosophical Psychology* 24: 252–58.
- KNOBE, J., and J. PRINZ. Forthcoming. Intuitions about consciousness: Experimental studies. *Phenomenology and Cognitive Science*.
- KNOBE, J., and E. ROEDDER. n.d. The concept of valuing: experimental studies. Unpublished manuscript.
- LESLIE, A., J. KNOBE, and A. COHEN. 2006. Acting intentionally and the side-effect effect: 'Theory of mind' and moral judgment. *Psychological Science* 17: 421–27.
- LEWIS, D. 1972. Psychophysical and theoretical identifications. *Australasian Journal of Philosophy* 50: 249–58.
- LEVY, N. n.d. Experimental philosophy: a critique. Unpublished manuscript.
- MACHERY, E. Forthcoming. Understanding the folk concept of intentional action: Philosophical and experimental issues. *Mind & Language*.
- MACHERY, E., R. MALLON, S. NICHOLS, and S. P. STICH. 2004. Semantics, cross-cultural style. *Cognition* 92 (3): 1–12.
- MACKIE, J. 1977. *Ethics: Inventing right and wrong*. London: Penguin.

- MALLE, B. 2001. Folk explanations and intentional action. In *Intentions and intentionality: Foundations of social cognition*, edited by B. F. Malle, L. J. Moses, and D. A. Baldwin. Cambridge, Mass.: MIT Press.
- . 2006. Intentionality, morality, and their relationship in human judgment. *Journal of Cognition and Culture* 6 (1/2): 87–112.
- MALLE, B., and J. KNOBE. 1997. The folk concept of intentional action. *Journal of Experimental Social Psychology* 33: 101–21.
- MALLE, B., and S. NELSON. 2003. Judging *mens rea*: The tension between folk concepts and legal concepts of intentionality. *Behavioral Sciences and the Law* 21: 563–80.
- MALLON, R., E. MACHERY, S. NICHOLS, and S. STICH. Forthcoming. Against arguments from reference. In *Metametaphysics*, edited by D. Chalmers, D. Manley, and R. Wasserman. Oxford: Oxford University Press.
- MCCANN, H. 2005. Intentional action and intending: Recent empirical studies. *Philosophical Psychology* 18 (6): 737–48.
- MEEKS, R. 2004. Unintentionally biasing the data: Reply to Knobe. *Journal of Theoretical and Philosophical Psychology* 24: 220–23.
- MELE, A. 2001. Acting intentionally: Probing folk intuitions. In *Intentions and intentionality: Foundations of social cognition*, edited by B. F. Malle, L. J. Moses, and D. A. Baldwin. Cambridge, Mass.: MIT Press.
- . 2003. Intentional action: controversies, data, and core hypotheses. *Philosophical Psychology* 16: 325–40.
- MYNATT, C. R., M. E. DOHERTY, and R. D. TWENEY. 1977. Confirmation bias in a simulated research environment: An experimental study of scientific inference. *Quarterly Journal of Experimental Psychology* 29: 85–95.
- NADELHOFFER, T. 2004a. The Butler problem revisited. *Analysis* 64 (3): 277–84.
- . 2004b. Praise, side effects, and intentional action. *Journal of Theoretical and Philosophical Psychology* 24: 196–213.
- . 2005. Skill, luck, control, and intentional action. *Philosophical Psychology* 18 (3): 343–54.
- . 2006a. Bad acts, blameworthy agents, and intentional actions: Some problems for jury impartiality. *Philosophical Explorations* 9 (2): 203–20.
- . 2006b. Foresight, moral considerations, and intentional actions. *Journal of Cognition and Culture* 6 (1): 133–58.
- NADELHOFFER, T., and A. FELTZ. Forthcoming. Folk intuitions, slippery slopes, and necessary fictions: An essay on Smilansky's free will illusionism. *Midwest Studies in Philosophy*.
- NADELHOFFER, T., and E. NAHMIAS. Forthcoming. Polling as a valuable pedagogical tool for teaching philosophy. *Teaching Philosophy*.
- NADELHOFFER, T., and E. NAHMIAS. n.d. Experimental philosophy: A house divided? Unpublished manuscript.
- NAHMIAS, E. 2006. Folk fears about freedom and responsibility: Determinism vs. reductionism. *Journal of Cognition and Culture* 6 (1/2): 215–37.
- NAHMIAS, E., D. COATES, and T. KVARAN. Forthcoming. Intuitions about free will and moral responsibility: Mapping the terrain. *Midwest Studies in Philosophy*.
- NAHMIAS, E., S. MORRIS, T. NADELHOFFER, and J. TURNER. 2004. The phenomenology of free will. *Journal of Consciousness Studies* 11 (7/8): 162–79.
- . 2005. Surveying freedom: Folk intuitions about free will and moral responsibility. *Philosophical Psychology* 18 (5): 561–84.
- . 2006. Is incompatibilism intuitive? *Philosophy and Phenomenological Research* LXXIII, No 1.

- NICHOLS, S. 2002. Norms with feeling: Towards a psychological account of moral judgment. *Cognition* 84: 221–36.
- . 2004a. Folk concepts and intuitions: From philosophy to cognitive science. *Trends in Cognitive Science* 8 (11): 514–18.
- . 2004b. After objectivity: An empirical study of moral judgment. *Philosophical Psychology* 17: 5–28.
- . 2004c. Folk psychology of free will. *Mind & Language* 19: 473–502.
- . 2006a. Folk intuitions on free will. *Journal of Cognition and Culture* 6 (1/2): 57–86.
- . 2006b. Free will and the folk: Response to commentators. *Journal of Cognition and Culture* 6 (1/2): 305–20.
- NICHOLS, S., and J. KNOBE. Forthcoming. Moral responsibility and determinism: The cognitive science of folk intuition. *Nous*.
- NICHOLS, S., and R. MALLON. 2006. Moral dilemmas and moral rules. *Cognition* 100: 530–42.
- NICHOLS, S., S. STICH, and J. WEINBERG. 2003. Metaskepticism: Meditations in ethno-epistemology. In *The skeptics*, edited by S. Luper. Burlington, Vt.: Ashgate.
- NICHOLS, S., and J. ULATOWSKI. Forthcoming. Intuitions and individual differences: The Knobe effect revisited. *Mind & Language*.
- NICKERSON, R. S. 1998. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology* 2: 175–220.
- NISBETT, R. E. 2003. *The geography of thought: How Asians and Westerners think differently . . . and why*. New York: Free Press.
- NISBETT, R. E., and L. ROSS. 1980. *Human inference: Strategies and shortcomings of social judgment*. Englewood Cliffs, N.J.: Prentice-Hall.
- PHELAN, M., and H. SARKISSIAN. Forthcoming. The folk strike back: Or, why you didn't do it intentionally, though it was bad and you knew it. *Philosophical Studies*.
- PIZZARRO, D., and P. BLOOM. 2003. The intelligence of moral intuitions: A comment on Haidt (2001). *Psychological Review* 110 (1): 193–96.
- PIZZARRO, D., E. UHLMANN, and P. BLOOM. 2003. Causal deviance and the attribution of moral responsibility. *Journal of Experimental Social Psychology* 39 (6): 653–60.
- PIZZARRO, D., E. UHLMANN, and P. SALOVEY. 2003. Asymmetry in judgments of moral blame and praise: The role of perceived metadesires. *Psychological Science* 14: 267–72.
- ROSS, L., D. GREENE, and P. HOUSE. 1977. The false consensus effect: an egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology* 13: 279–301.
- SARDJEVLADŽE, L., and E. MACHERY. n.d. Trusting and punishing artifacts. Unpublished manuscript.
- SMITH, M. 1994. *The moral problem*. Cambridge: Basil Blackwell.
- SOSA, E. 2005. A defense of the use of the intuitions in philosophy. In *Stich and his critics*, edited by M. Bishop and D. Murphy. Oxford: Blackwell Publishers.
- . 2007. Experimental philosophy and philosophical intuition. *Philosophical Studies* 132 (1): 99–107.
- STICH, S. 1983. *From folk psychology to cognitive science: The case against belief*. Cambridge, Mass.: MIT Press.
- . 1984. Relativism, rationality and the limits of intentional description. *Pacific Philosophical Quarterly* 65 (3): 211–35.
- . 1985. Could man be an irrational animal? *Synthese* 64 (1): 115–35.
- . 1990. *The fragmentation of reason*. Cambridge, Mass.: MIT Press.
- . 1994. Philosophy and psychology. In *A companion to the philosophy of mind*, edited by S. Guttenplan. Cambridge: Basil Blackwell.

- . 1996. *Deconstructing the mind*. Oxford: Oxford University Press.
- . 1998. Reflective equilibrium, analytic epistemology, and the problem of cognitive diversity. In *Rethinking intuition: The psychology of intuition and its role in philosophical inquiry*, edited by M. DePaul and W. Ramsey. Lanham, Md.: Rowman and Littlefield.
- STICH, S., and R. NISBETT. 1980. Justification and the psychology of human reasoning. *Philosophy of Science* 47 (2): 188–202.
- STICH, S., and J. WEINBERG. 2001. Jackson's empirical assumptions. *Philosophy and Phenomenological Research* 62: 637–43.
- STOTZ, K. Forthcoming. How to test philosophical analyses of a scientific concept: Notes from the field. *Philosophy of Science*.
- STOTZ, K., and P. E. GRIFFITHS. 2004. Genes: Philosophical analyses put to the test. *History and Philosophy of the Life Sciences* 26 (1): 5–28.
- STOTZ, K., P. E. GRIFFITHS, and R. D. KNIGHT. 2004. How scientists conceptualize genes: An empirical study. *Studies in History and Philosophy of Biological and Biomedical Sciences* 35 (4): 647–73.
- STRAWSON, G. 1994. The impossibility of moral responsibility. *Philosophical Studies* 75: 5–24.
- SWAIN, S., J. ALEXANDER, and J. WEINBERG. Forthcoming. The instability of philosophical intuitions: Running hot and cold on Truetemp.
- TURNER, J., and E. NAHMIA. 2006. Are the folk agent causationists? *Mind & Language* 21 (5): 597–609.
- VARGAS, M. 2005. The revisionist's guide to responsibility. *Philosophical Studies* 125 (3): 399–429.
- VOHS, D., and J. SCHOOLER. n.d. The price of determinism: Denying the existence of free will increases cheating behavior. Unpublished manuscript.
- WASON, P. C. 1960. On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology* 12: 129–40.
- WEGNER, D. 2002. *The illusion of conscious will*. Cambridge, Mass.: MIT Press.
- WEINBERG, J., S. NICHOLS, and S. STICH. 2001. Normativity and epistemic intuitions. *Philosophical Topics* 29 (1/2): 429–60.
- WILLIAMSON, T. 2004. Philosophical 'intuitions' and scepticism about judgment. *Dialectica: International Journal of Philosophy of Knowledge* 58 (1): 108–53.
- WOOLFOLK, R., and J. DORIS. 2002. Rationing mental health care: Parity, disparity, and justice. *Bioethics* 16: 469–85.
- WOOLFOLK, R., J. DORIS, and J. DARLEY. Forthcoming. Attribution and alternate possibilities: Identification and situational constraint as factors in moral cognition. *Cognition*.
- YOUNG, L., F. CUSHMAN, R. ADOLPHS, D. TRANEL, and M. HAUSER. 2006. Does emotion mediate the relationship between an action's moral status and its intentional status? Neuropsychological evidence. *Journal of Cognition and Culture* 6 (1/2): 291–304.

Thomas Nadelhoffer (author to whom correspondence should be addressed), Dickinson College, Department of Philosophy, Carlisle, PA 17013, USA. E-mail: tnadelhoffer@gmail.com

Eddy Nahmias, Brains & Behavior Program, Department of Philosophy, Georgia State University, 1116 One Park Tower, Atlanta, GA 30302-4089, USA. E-mail: enahmias@gsu.edu