

Experimental Philosophy and the Fruitfulness of Normative Concepts

Matthew Lindauer
Brooklyn College, City University of New York

Forthcoming in *Philosophical Studies*

Abstract: This paper provides a new argument for the relevance of empirical research to moral and political philosophy and a novel defense of the positive program in experimental philosophy. The argument centers on the idea that normative concepts used in moral and political philosophy can be evaluated in terms of their fruitfulness in solving practical problems. Empirical research conducted with an eye to the practical problems that are relevant to particular concepts can provide evidence of their fruitfulness along a number of dimensions. An upshot of the argument is that philosophers should not only engage with but must also be involved in conducting experimental studies that examine the practical roles that normative concepts can play. Rather than just clearing the way for philosophical work to be done, the argument has the further implication that empirical research will be required to advance at least some important debates in moral and political philosophy.

Proponents and critics of experimental philosophy have distinguished between the field's "negative" and "positive" programs.¹ The negative program has so far attracted the most attention, both in the academy and in the popular press.² This is perhaps not surprising, given the radical critiques of a priori philosophical methods that early papers in the negative program put forward (Weinberg 2007; Alexander and Weinberg 2007; Alexander et al. 2014). Some critics

¹ Nadelhoffer and Nahmias (2007); Weinberg (2007); Alexander and Weinberg (2007); Cappelen (2012); Alexander et al. (2014); Fisher (2015); Shepherd and Justus (2015); Williamson (2016). Knobe (2016) argues that the positive-negative program distinction fails to capture much of the important work being conducted in the field. Sytsma and Livengood (2016) propose a broader taxonomy of programs (non-intuitional, negative, positive, cognitive, and descriptive) and hold that the work that Knobe points to largely falls under what they refer to as the "cognitive program."

² See, e.g., the discussion in the Opinion Pages of the New York Times ("X-Phi's New Take on Old Problems" 2010).

of experimental philosophy direct their objections almost exclusively at the negative program, and seem to concede some ground to what could be thought of as the positive program (Williamson 2016). Yet the positive program has not been given a robust explanation or defense.³ If all experimental philosophy falls into either the negative or positive program, it could be that many have been inclined to think of the positive program merely in contrast to the more widely discussed negative program.

Joshua Knobe has argued that the positive/negative program distinction, cashed out in terms of making a positive contribution to conceptual analysis or providing evidence against the methodological assumptions of conceptual analysis, is not helpful (Knobe 2016). A broader distinction, however, can be drawn between experimental work that views itself as continuous with traditional a priori philosophical methods and that which views itself as discontinuous with or undermining those methods. The positive program, if understood in this way, encompasses much of the work that Knobe points to in showing that there are overwhelming instances of research in experimental philosophy that fall outside of what could be regarded as the negative program (or the positive program, on the understandings of these terms that he is using). Many philosophers and psychologists want to better understand how empirical work contributes to philosophical questions, and so the notion that such empirical work is continuous with traditional philosophical methods requires justification.

This paper provides an argument in support of the positive program, understood as the program of bringing empirical research to bear on philosophical questions in ways that are continuous with traditional a priori research methods. I argue that philosophers should be involved in conducting at least some of the empirical research that bears on issues in moral and

³ Of course, there are some exceptions. See, e.g., Shepherd and Justus (2015).

political philosophy.⁴ Many theorists hold the view that the normative concepts employed in these fields,⁵ in addition to being coherent and well motivated from a purely theoretical perspective, should be evaluated in part by the practical role that they play. The extent to which moral and political concepts play their practical role well I will call “fruitfulness,” a term that I borrow from Carnap (1950).⁶ The fruitfulness of moral and political concepts, on my view – and by concepts, I intend to capture also the norms, theories, and principles that fall under them, or “conceptions”⁷ of concepts – can only be determined with the help of empirical investigation. In particular, determining whether or not a given moral or political concept will be fruitful in solving practical *problems* is a question that requires empirical knowledge to answer. The argument therefore provides a new justification for the “positive program” in experimental philosophy in terms of its importance to moral and political philosophy. Indeed, my view has the implication that the positive program is not merely in the service of philosophical research, but deeply continuous with the goals and tendencies of philosophical research in these areas.

⁴ While I focus primarily on psychological research, the argument that I present in this paper supports a wide role for empirical research from other social sciences in moral and political philosophy. It may also support a broader role for non-empirical research in some of these fields, including economic theory, social choice theory, and game theory, but establishing this point would go beyond the bounds of the present paper. I am grateful to an anonymous reviewer for suggesting that my argument may also encompass some non-empirical research.

⁵ I will remain neutral on the degree to which normative concepts employed in other fields of philosophy, such as epistemology or aesthetics, can be similarly evaluated in terms of their fruitfulness. Some commentators have suggested that what I have to say here about moral and political concepts can be extended to these other fields.

⁶ Carnap used this term in putting forward a somewhat different view of the practical role that scientific concepts play. I discuss this in Section 4.

⁷ The distinction between concepts and conceptions is due to Rawls (1971).

1. The Fruitfulness of Normative Concepts

There's an unspoken agreement between positive program experimental philosophers and moral and political philosophers who are receptive to experimental philosophy but do not conduct empirical research. The former won't press on where the intuitions come from or the scientific explanations of where our judgments get their hold if the latter don't make them say why the social sciences are important to ethics. Instead of leaving this situation as it stands, we should be confident that good explanations can be given on the second issue that are compatible with non-reductive, non-skeptical accounts of the first issue. This paper provides one such attempt at addressing the second issue, by showing that in order to evaluate the fruitfulness of moral and political concepts, research in the social sciences must be engaged with and, in my view, even conducted by philosophers.

One way into thinking about the fruitfulness of these normative concepts is to consider the purposes for which we engage in moral inquiry. Aristotle famously held that human beings have a *telos* or an end, but his remarks on the end of ethical study are often overlooked. In *Nicomachean Ethics*, Book II, he claims that we study ethics for a different reason than we study other areas of philosophy and subject matters: "we are not conducting this inquiry in order to know what virtue is, but in order to *become good*" (Aristotle 1999, Book II.2, emphasis mine). The study of ethics has a point for Aristotle, to help us figure out how to live our lives in the best way, which for him is the life of virtue. I mention this remark not because I wish to defend Aristotelian virtue ethics, but because we can see what is deeply correct about Aristotle's general claim while expanding the range of purposes that moral and political concepts might have.

Normative concepts help us to live together with other human beings by aiding us in coordinating our behavior and cooperating with one another (Kitcher 2014). Working in moral and political philosophy, we are developing on a pre-existing tapestry of norms, theories, and principles, which sometimes cluster around a particular concept, and that have helped us to avoid conflict and inefficiency better than we would have without them.⁸ As I will demonstrate, moral and political concepts can be and often are evaluated in terms of their “fruitfulness,” by which I mean how well they help us to solve practical problems, problems that we inevitably face as human beings. Many moral and political philosophers are committed to the fruitfulness of normative concepts mattering to their evaluation, and in this minimal sense are sympathetic to the Aristotelian understanding of the purposes of moral inquiry.⁹

Of course, there is another tradition in moral and political philosophy that eschews fruitfulness, aligning more closely with Plato’s views than those of Aristotle that I have just mentioned. According to this line of normative inquiry, moral and political concepts are no different than non-normative concepts, in the sense that the practical roles that they might play should be extrinsic to any evaluations we might make of them.¹⁰ Norms, theories, and principles that fall under these concepts are to be examined in terms of theoretical virtues that exclude any role that these concepts have in our practical deliberations or solving problems that arise in our interactions with one another.

⁸ Garner (2007) disagrees, but the view that morality should be abolished is not widely held. Walker (2007) emphasizes the negative social roles that moral and political concepts often play, but doesn’t take this to show that we should try to do away with them.

⁹ This commitment does not require one to adopt any substantive positions in moral or political philosophy that Aristotle or Aristotelians endorse. It is instead a general commitment regarding how normative concepts should be evaluated, where their fruitfulness should go into such evaluations.

¹⁰ For recent defenses of this view, see Cohen (2008), Enoch (2011), and Estlund (2014).

Fully examining and adjudicating the debate between these long-standing camps in moral and political philosophy is beyond the bounds of this paper.¹¹ Rather than doing so, my goal in this paper is more modest. My primary task is to illustrate the various ways empirical research can be brought to bear on the fruitfulness of normative concepts. And in the course of doing so, I will point out the many instances in which moral and political philosophers have committed themselves, either explicitly or implicitly, to the fruitfulness of these concepts and the norms, theories, and principles that fall under them mattering for their evaluation. Rather than directly arguing against the Platonic view or for the Aristotelian view, then, the paper will make salient many of the types of research that bear on the kinds of questions that the Aristotelian view draws our attention to. And further, it will make clear what is at stake in adopting the Aristotelian view, what is made on limits in the study of normative concepts in doing so, and what cannot be part of the examination of normative concepts if we separate their evaluation from their fruitfulness.

I will therefore begin by pointing out a number of ways in which moral and political concepts can be and often are evaluated in terms of their fruitfulness. I will often refer to norms, theories, and principles that fall under a concept, what John Rawls (1971/rev. 1999) referred to as “conceptions” of concepts, because they are typically the items that are evaluated using practical criteria.

Starting with the view that normative concepts are fruitful to the extent that they help us to solve practical problems, a number of roles can be seen for these concepts and the norms,

¹¹ See Korsgaard (2003) for a characterization of this disagreement in terms of the history of the debate between realism and constructivism in moral philosophy. One need not be a constructivist to endorse what I will say here, but I will suggest that Korsgaard is committed to the fruitfulness of normative concepts mattering to their evaluation in Section 2.

theories, and principles falling under them in addressing practical problems. To understand which normative concepts can help us to solve the practical problems that they are supposed to, notably, empirical facts come into play. These facts include: 1) how people will act when particular norms, theories, and principles are internalized (call this “Motivation”); 2) how people will be prevented from acting when such norms, theories, and principles are internalized (“Prevention”); 3) whether when norms, theories, and principles compete with bias and other problematic such items, some are more effective in “fighting back” (“Resilience”); 4) whether some norms, theories, and principles are better subjects of consensus among people with other commitments that diverge but are nonetheless consistent with good will and a desire for peaceful coexistence (“Consensus”); and 5) whether these norms, theories, and principles provide or help us to formulate more useful prescriptions than others with respect to problems that we must solve (“Guidance”). Each of these roles highlight that, at least in the normative domain, empirical facts can help us to see whether concepts can help us to solve important practical problems:

1) How people will act when such norms, theories, and principles are internalized (Motivation)

It is not uncommon for philosophers to be concerned with evaluating moral and political concepts in terms of the effects that would result from their adoption. In examining these effects, one natural question to ask is how people will act when they employ these concepts and the norms, theories, and principles that fall under them. The self-effacingness objection to consequentialism is one example of a challenge to the fruitfulness of a moral theory in terms of how the theory would motivate people to act. According to that objection, it is a strike against consequentialist moral theories that they recommend that agents not internalize

consequentialism, in the sense of using a consequentialist decision procedure for action. Agents that did so would not be motivated to perform actions that meet the consequentialist's own requirement to maximize good consequences, and by this requirement it would be better, the objection continues, if they followed their moral intuitions¹² or some other non-consequentialist guide to action. This is the sense in which consequentialism is supposed to be "self-effacing." I am not asserting that the self-effacingness objection to consequentialism is successful. Many responses have been offered to it (see, e.g., Parfit 1992). My point is that this well-known objection to consequentialism suggests that philosophers are interested in what I am referring to as the *motivational* fruitfulness of a normative concept – what human beings who have internalized the concept or the norms, theories, and principles that fall under it, will be motivated to do.

Another example of a debate where the motivational force of normative concepts has been thought to matter is found in the literature on duties to the global poor. Peter Singer (1972) famously argued that we can extract stringent moral duties to help people in severe poverty in developing countries from principles of beneficence that we are committed to. Alison Jaggar (2001, 2005) and Thomas Pogge (2008) have argued, however, that while principles of beneficence also apply, more stringent duties against committing harm following from principles of justice have been violated by the affluent that make the requirement to address global poverty even stronger. Many have claimed that there will be a motivational difference between arguments that appeal to these two different types of principles because the principles of justice in question involve wrongful actions that have been committed, namely the collective imposition

¹² The assumption here, of course, is that they have the kinds of deontological moral intuitions that most people report. If the person in question has consequentialist moral intuitions, following those intuitions will involve acting on consequentialist principles.

of an unjust global order and a history of violent oppression through colonialism and imperialism. These negative duties are thought by some to be more motivating than the positive duties that Singer points to, which only focus on our failing to help the global poor, morally wrong *omissions* rather than actions. It has been taken as a virtue of the negative duty arguments that they will be more motivating than the positive duty argument with respect to global poverty (Pogge 2008; Lichtenberg 2010, 2014; Lawford-Smith 2012), at least when they are accepted. On the other hand, negative duty arguments may not succeed in being more motivating because people may not find it plausible that they are harming the global poor. This may be the case in part because they wish to avoid feeling guilty about contributing to severe poverty, whereas it is easier to accept the claim that we are merely failing to help (Lichtenberg 2010). I am again not taking a position in this particular debate but rather pointing to another example where the motivational fruitfulness of philosophical concepts, in this case principles of beneficence and justice as applied to a particular real-world moral problem, is thought to be relevant to their evaluation.

Empirical studies looking at the attitudes and behaviors of individuals in response to presenting or eliciting normative concepts can bring out the tendencies those normative concepts produce in action. This can be thought of as adopting an “intervention-based method” for the study of the motivational fruitfulness of these concepts. Linda Skitka’s research suggests that moral attitudes are especially strong and motivating commitments (Skitka et al. 2005; Skitka 2010), but little work has been done thus far to directly examine whether and when moral concepts and arguments can be used to motivate moral behavior. In the case of the self-effacingness objection to consequentialism, the intervention-based method might involve priming participants to think about a given problem in consequentialist terms and seeing if their

actions in the experiment maximize some measure of good consequences. Another example of the intervention-based method would involve presenting the arguments that Singer, Jagger, and Pogge have made regarding global poverty to different groups of participants and seeing both whether participants find them plausible and, among those who do find them plausible overall, how motivated they are to adopt actions to address global poverty, such as donating money that they are paid for completing the study to an effective organization working to alleviate it.¹³ A non-intervention based method, by contrast, would draw out the motivational consequences of pre-existing acceptance and internalization of particular normative items in different individuals and communities. Using this second method, we might look at the prior tendencies of different people towards more consequentialist or more deontological modes of ethical reasoning and how these people respond to decision-making tasks assigned to them. With respect to global poverty, we can study the donation behavior of participants who differ in whether they view global poverty primarily in terms of duties of beneficence or justice and compare the motivational force of these ideas in this way, controlling for confounding factors such as political orientation, level of income, religiosity, and so on.

These are just some of the types of empirical studies that may be conducted to examine the motivational force of normative concepts. Other studies might examine the comparative or combined effects of normative concepts and appeals to emotion. Deborah Small, George Loewenstein, and Paul Slovic (2007) famously showed that the photo of a child (“Rokia”) living in poverty garnered significantly higher charitable donations than that photo accompanied by statistical information about other children in the same situation. Yet emotional appeals also are

¹³ Luke Buckland, David Rodríguez-Arias, Carissa Véliz, and I have conducted empirical studies examining the effects of reading some of these arguments on charitable giving. See Buckland et al. (unpublished manuscript).

limited in their effectiveness by “psychic numbing,” our tendency to be able to focus our concern on only one individual, even when the appeal is made to our emotions rather than our rational capacities (Slovic 2007). From these two points, Slovic concludes that we must seek institutional solutions to pressing moral problems like global poverty and the general disinterest with which rich nations have approached genocide in other countries (Slovic 2007, p. 91). One might also conclude, however, that we should be examining which moral norms, theories, and principles do well in combining with appeals to our emotions to generate the greatest motivational effects.¹⁴

Matthew Feinberg’s research suggests that if you want to motivate people with arguments, these arguments should be framed in terms of moral values that they already hold (Feinberg and Willer 2015). It may then be that the motivational fruitfulness of a moral argument for a broad population depends, at least in part, on the presentation of the argument making reference to shared values, such as the values of harm and care (Graham et al. 2011; Haidt 2012). Relatedly, one might examine whether different arguments work particularly well in specific populations and not others. Fiery Cushman’s research on model-free versus model-based moral reasoning (Cushman 2013) also suggests that more deontological and more consequentialist moral arguments may have relative costs and benefits in different contexts. By assigning a moral evaluation to action types rather than examining the consequences of each action directly, deontological frameworks may be computationally light, yielding lesser burdens on people already facing limited cognitive resources and time to devote to moral causes (*ibid.*, 278). On the other hand, consequentialist moral frameworks may be more flexible and allow for

¹⁴ Peter Singer, Paul Slovic, Daniel Västfjäll, Joshua Greene, Marcus Mayorga and I have recently been conducting empirical research on these issues.

the reexamination of the way particular actions are evaluated in the context of new moral problems that demand reconceptualization to motivate people (*ibid.*, 277-278), such as the problem of climate change. Each of these lines of research should inform intervention and non-intervention based approaches that seek to go beyond mere speculation and actually examine the effects that moral concepts and arguments have on moral motivation.

Nicole Hassoun, Emir Malikov, and Nathan Lubchenco (2016) have also taken up the non-intervention based method of studying the principles that seem to best capture people's donation behavior in funding microloans to persons in developing countries. Non-intervention based approaches may serve as a way of testing the extensional adequacy of empirical claims regarding what kinds of moral attitudes and principles tend to motivate particular moral behaviors. For instance, we might examine the extent to which individuals' moral foundations (Graham et al. 2011; Haidt 2012; Feinberg and Willer 2015) or sense of their own moral identity (Bryan et al. 2011, 2012) predict which moral principles provide the best unifying explanations for their actions when they are given an opportunity to address a moral problem like global poverty. These and other examples may be furnished as possible ways of examining the motivational effects of normative concepts when they have been internalized and assessing normative concepts, in part, on this basis.

2) How people will be prevented from acting when such norms, theories, and principles are internalized (Prevention)

In addition to intervention and non-intervention based approaches to looking at the effects of normative concepts in promoting certain types of behavior, empirical research can be

useful in assessing the ability of normative concepts to *prevent* harmful behaviors. Jeanette Kennett and Cordelia Fine (2009) survey the literature suggesting that moral judgments are not merely post hoc rationalizations (contra Haidt 2001) but can help us to regulate our behaviors prospectively.¹⁵ They convincingly argue that judgments about moral principles can have prosocial effects when used in the right ways. Consider Daniel Batson's important research (see, e.g., Batson 2008) showing that people's tendencies to cheat and play unfairly in games can be inhibited by making them feel that they are being observed. Notably, this awareness doesn't have to be conscious and may be activated by even as minimal a stimulus as a mirror or representation of a pair of eyes (see also Bateson et al. 2006 and Ernest-Jones et al. 2011). Empirical research could study, for instance, whether there is a difference in the ability to suppress the tendency to cheat or act selfishly in individuals who employ either more consequentialist or deontological reasoning. Robert Wicklund (1979) also suggests that promoting objective self-awareness leads to better behavior, whereas Kathleen Vohs and Jonathan Schooler (2008) find that a free-will manipulation increasing subject's sense that they have free will improves their conduct. It may then be that moral theories that emphasize these notions, such as Kantian deontology, which focuses heavily on the notion of a singular free rational subject, will have the tendency to produce better behavior than moral theories issued from the "point of view of the universe" such as consequentialism (Sidgwick 1907). The fruitfulness of normative concepts can be assessed not only by what they directly promote, but also what they prevent in terms of human tendencies to act in ways that are unfair or harmful to others.

¹⁵ See, e.g., Barrett et al. (2004), Payne (2005), and Amodio et al. (2008).

3) When some norms, theories, and principles compete with bias and other problematic such items, which are more effective in “fighting back” (Resilience)

Related to the prior point, there is the matter of pre-existing biases and attitudes that are problematic and tend to influence behavior but do not always do so. Implicit biases and stereotype threat can affect people who do not explicitly endorse the relevant attitudes.¹⁶ Jennifer Lerner and Philip Tetlock (1999) have found that knowing that we will be accountable to others for explaining our actions typically inhibits bias and improves judgment. Kennett and Fine (2009) discuss Keith Payne’s research suggesting that we can inhibit the behavioural expression of bias through goal-directed processing, such as when one has a salient egalitarian goal of preventing oneself from stereotyping others (Payne 2005). Betsy Levy Paluck’s research (Levy Paluck et al. 2016) also suggests that with certain practical problems, such as bullying, theories that make reference to individuals who are viewed as sources of normative information may be more motivating than theories that are presented in more general terms. This research may provide reason to think that virtue theoretic approaches do, contrary to common criticisms, have distinctive practical benefits. Empirical methods can be used to assess whether certain normative concepts or theoretical approaches are better or worse in empowering people to resist the influence of implicit biases and other problematic attitudes across a range of issues.

¹⁶ Recent studies have questioned the extent to which both implicit bias and stereotype threat affect the kinds of behaviors thought to be affected by them in prior research. While some of these effects have not replicated in these studies, other effects, such as the “shooter bias” effect (Correll et al. 2002), have replicated (Essien et al. 2017). There is an ongoing debate about the effects of implicit bias and stereotype threat on particular behaviors and what conclusions can be drawn from this literature for public policy. The point that I am making here relies only on the claim that any negative effects of biases grounded in social group membership on evaluations of others and oneself are worth addressing through the use of normative concepts if doing so is possible, and there is some evidence that this can be done (see Payne 2005).

Sally Haslanger (2000, 2012) focuses on this potential effect of a concept in putting forward a feminist, antiracist theory of the concepts of race and gender. “At the most general level,” she writes, “the task is to develop accounts of gender and race that will be effective tools in the fight against injustice” (2000, p. 36). I will discuss Haslanger’s view and how what I’m calling “resilience” relates to it in Section 5.

4) Whether some norms, theories, and principles are better subjects of consensus among people with other commitments that are consistent with good will and a desire for peaceful coexistence (Consensus)

A further criterion of fruitfulness for normative concepts concerns whether or not such concepts are better able to secure consensus among people with other commitments that differ and who are at least not opposed to cooperation and coexistence. An example of the sort of research that examines the degree to which a normative concept is fruitful in this way can be found in the empirical literature on Rawls’ theory of justice. Norman Frohlich and Joe Oppenheimer (1987) were some of the first social scientists to look at whether people placed in a situation that is meant to model Rawls’ original position would select principles of justice of the kind that Rawls predicted. Strikingly, they found that their subjects in this setup tended to endorse a restricted utilitarian principle, which allows for the maximization of utility once everyone meets a basic minimum standard, rather than the difference principle. Because this research focuses on income distributions, and not primary goods in general, their result is not strictly incompatible with Rawls’ prediction. Still, this experiment provides at least suggestive evidence that a restricted utilitarian principle will achieve consensus among subjects in the

original position, whereas Rawls holds that the difference principle will defeat this principle (2001, pp. 126–130). Their finding has been replicated (Lissowski et al. 1991; Bruner 2018). However, some have argued against the idea that we will be able to find either a single, uniquely choiceworthy set of principles of justice (Sugden 1990; Skyrms 1996, 2016; Thrasher 2013; Muldoon et al. 2014) or procedure for arriving at such principles (Gaus 2010; Moehler 2018; Bruner and Lindauer forthcoming). Such criticism is compatible, of course, with the goal of finding principles and procedures that will be widely accepted. At the level of individual morality, Mikhail (2007; 2011) has sought to uncover a “universal moral grammar” that underlies the ethical frameworks that all cultures possess. This fundamental layer of moral cognition may be used to derive clues regarding which normative concepts will tend to secure the most widespread agreement. Empirical methods can shed light on both the uptake¹⁷ of normative concepts in varied individuals and how well these concepts provide a common ground for working together to solve practical problems.

5) Whether norms, theories, and principles provide or help us to formulate more useful prescriptions than others with respect to problems that we must solve (Guidance)

¹⁷ Working largely within the framework of Carnapian explication, Pinder (forthcoming) argues that experimental philosophy can help to determine whether explicated concepts are fruitful in the sense that they can achieve uptake by “the relevant theoretical community.” Although I find much to agree with in Pinder, my view is not offered within the Carnapian framework, with its emphasis on explication. It is also a commitment of my view that, rather than only achieving uptake by moral and political philosophers, ordinary people as moral agents and citizens are part of the relevant community that moral and political philosophy must be fruitful for. This commitment, I suggest throughout the paper, is not uncommon among moral and political philosophers who take the practical upshots of their views to be philosophically significant.

Many philosophers hold that a moral theory should be action guiding, or at least that being action-guiding is a strike in a moral theory's favor. For instance, a traditional objection to Aristotelian virtue ethics is that it doesn't provide us with sufficient instructions for how to respond morally to situations that we find ourselves in, merely telling us to do as the virtuous person would do. I have suggested above that I think this objection ignores ways in which virtue ethical approaches may have distinctive practical benefits, but it is a common objection to such approaches and one that presupposes that moral theories should be useful in guiding action. Philosophers have also recently been interested in how feasibility constraints should enter into our moral and political theorizing, and whether a theory takes such constraints into account affects its action-guiding potential. As in ethics, one position in political philosophy involves the view that feasibility constraints are irrelevant to our theorizing about justice. But other philosophers have thought that, at least with some of our purposes as moral and political philosophers, feasibility constraints should be taken into account (Southwood 2016). Part of why we should care about feasibility constraints, these philosophers think, is because we need moral and political philosophy to issue us guidance about what to do about real world moral problems. This is often how the distinction between ideal and non-ideal theory is cashed out, where non-ideal theory is meant to play a role in helping us to address such problems (see, e.g., Valentini 2012).¹⁸ Some philosophers have argued that their theories of political justice (Rawls 1971/rev. 1999, 1980, 1993) or freedom (Pettit 1999, 2014) provide better guidance than their alternatives and that this is a mark in their favor. For Rawls, the "task" of political philosophy is "to articulate a public conception of justice that all can live with who regard their person and their relation to society in a certain way" (Rawls 1980, p. 519). He goes on to make the even

¹⁸ Mills (2005). See Valentini (2012) for a helpful overview of the ideal/non-ideal theory debate.

stronger claim that while “doing this may involve settling theoretical difficulties, the practical social task is primary.” In addition to being the subject of an “overlapping consensus” (Rawls 1971, 1993), a Rawlsian public conception of justice is supposed to guide citizens in developing and reforming their society’s political institutions on grounds that all citizens can accept. Philip Pettit refers to his republican conception of freedom as non-domination as a “moral compass” (2014) and describes it as a regulative ideal for assessing and reforming societal policies and institutions. He holds that from a purely theoretical perspective, it is less clear that we should prefer this conception to the alternative libertarian and liberal conceptions of freedom. But once we account for its superior ability to provide guidance on how to evaluate and reform a society’s institutions in line with the demands of justice, republican freedom is the winner among these competing conceptions of freedom.

Whether a given moral or political norm, theory, or principle provides helpful guidance concerning how to solve practical problems is an empirical fact. Studies can be designed, for instance, to see whether reflecting on people who are moral exemplars tends to lead subjects to behave more fairly than they otherwise would. There is evidence to suggest that thinking about positive exemplars from one’s own social group(s) can inhibit the effects of stereotype threat (Marx and Roman 2002; McIntyre et al. 2003, 2005). Further, it may also be the case that thinking about virtuous individuals, while this may not yield a determinate answer to what to do in scenarios like those presented in the trolley problem, will be more helpful than using the abstract, general principles of other approaches to moral theory in a variety of situations (Levy Paluck et al. 2016). If this were the case, as I have suggested above, it would imply that virtue ethics can be action guiding and perhaps have distinctive practical benefits. Similarly, we may test whether the Rawlsian conception of justice helps participants to divide resources up in ways

that other participants regard as fair, and whether Pettit's conception of freedom does allow persons to form consistent and plausible judgments about how to address societal injustices.

Other examples where political concepts have been thought to be more or less action guiding than alternatives are found in the literatures on racial inequality and on measures of welfare. In the literature on race, it has been argued that colorblind policies do not provide good guidance on how to solve problems of racial prejudice and inequality in societies with histories of racial injustice (Boxill 1992; Bonilla-Silva 2003; Anderson 2010; Mills 2011). In the literature on welfare, it has been argued that subjective measures won't guide us to improve the welfare of the worst off because their subjective reports are often influenced by the fact that marginalized people can have harmful adaptive preferences that are the result of internalizing their oppression (Khader 2011). Further empirical study will be useful in determining which moral and political concepts best respond to these problems.

The above examples are meant to give a better sense of how empirical research on these topics either has been or could be conducted. Of course, no one experiment or set of experiments is likely to settle the question of whether a given normative concept is, in the intended sense, a fruitful one. This is a caricature of the positive program. Ethics is not reduced to opinion polling by my approach, but rather takes into account empirical data that bears on the fruitfulness of the normative concepts at hand. Notably, there will typically be a gap between the data that we need in order to know which concepts can help us solve practical problems and the data available in a context where philosophers are not involved in conducting empirical research. I will revisit this point in Section 6.

2. Relations to Other Philosophical Views and Positions

So who would disagree with the view that I've put forward here? No philosophical view appeals to all audiences, and it would be surprising if this one were different. As noted above, there are views in moral and political philosophy that align more closely with Plato's views than the one that Aristotle puts forward in suggesting that ethical study has a practical point. In moral philosophy, hardcore or "robust" moral realists hold the view that facts about solving practical problems are strictly irrelevant to the evaluation of moral norms, theories, and principles (Enoch 2011). In political philosophy, the view is sometimes expressed that feasibility considerations and other practical facts should not enter into our evaluations of theories of political justice (Cohen 2008; Estlund 2014). I noted above that my goal in this paper is not to argue directly against these views, but rather to show the important implications for the relationship between experimental philosophy and the study of normative concepts that follow from an opposing view. At this point, I will expand the list of prominent philosophers and positions in philosophy that have been and should be receptive to what I have said here.

While I have already mentioned Aristotle's remarks about the practical importance of ethical study, it may be surprising that Kantian approaches to moral thinking, at least in the constructivist tradition, also provide an important line of support for thinking about the role of empirical research in ethics. Kant himself, of course, had some fairly dismissive things to say about moral philosophy that is not purely a priori (Kant 2012, Preface). Joshua Greene's research also purports to undermine Kantian deontology and support consequentialism (Greene et al. 2001; Singer 2005; Greene 2007), and this might be another cause for surprise. Yet there is reason to think that the Kantian constructivist tradition can provide some very helpful insights on

the relationship between moral thinking and the solution of practical problems. Christine Korsgaard (2003) examines a distinctive trend in moral thinking that emphasizes the relationship between ethics and practical reasoning. She argues that the history of constructivist and realist moral thought is a debate over whether or not, as the constructivist argues, a substantive account of morality grounded in practical reason is appropriate or whether, as realists tend to hold, constructivism doesn't give morality an independent enough status. While I don't view realists per se as committed to the view that moral concepts cannot be evaluated in terms of fruitfulness, it will be helpful to pause on Korsgaard's presentation of constructivism here for her useful formulation of the idea that moral concepts solve practical problems.

As I just mentioned, Korsgaard regards the substantive aspect of morality as that of practicality. In particular, she holds that constructivist views regard normative concepts as solutions to practical problems. For Kant, it is the problem of practical reason, a will needing a principle in order to be free.¹⁹ But Korsgaard notes that for Rawls, too, normative concepts can be evaluated in terms of whether they help us to solve practical problems. For Rawls, justice is a normative concept that is supposed to solve problems in liberal societies, societies where people differ in their moral, religious, and other personal commitments but are also committed to peaceful coexistence and mutual toleration (see especially Rawls 1993). Perhaps most obvious in debates surrounding Rawls' earlier work, which was less tied to the notion of a liberal society in particular, is the problem of distributive justice (Rawls 1971). Justice as fairness is a "conception" or specification of the concept of justice intended to solve that problem by offering

¹⁹ Street (2008) takes Korsgaard to task on this point, arguing that it is implausible to regard this as a practical problem. However, Street in her recent work suggests that metaethics may be therapeutic, helping us to deal with "the problem of attachment and loss" (Street 2016), and so it seems that she is receptive to the thought that moral theory could be evaluated in terms of serving a practical purpose.

an account of distributive justice, that, through the original position thought experiment, all should realize is fair to all. My view does not presuppose the truth of constructivism, as I've noted above. The point is that constructivists, and Korsgaard in particular, have given a particularly helpful characterization of the notion that normative concepts are meant to solve practical problems (for Korsgaard, they just *are* solutions to such problems).

There are a number of other issues and positions in moral and political philosophy that bear an important relation to practical action. Take the debate between internalists and externalists in ethics, for instance. For moral internalists, motivation is *internal* to moral judgment – there is a necessary connection between sincere moral judgments and motivation.²⁰ If someone sincerely judges that they ought to perform an action, according to this view, they must be at least somewhat motivated to perform that action. If they don't have any motivation to perform that action, in other words, the internalist denies that they sincerely judge that they are required to act in that way. This prominent brand of moral internalism is thus committed to a conceptual entailment between moral judgment and action.

Additionally, discussions of the ought-implies-can principle in moral and political philosophy show that many philosophers are committed to an important role for empirical facts in normative philosophical research. Many philosophers embrace this principle or hold that some nearby principle is correct (Southwood 2016). Most of these principles suggest that the set of our morally obligatory actions consists only in actions that we are capable of performing.

Here again, the link between practical action and morality is important, if in a different way. In

²⁰ For a helpful discussion of this view, which is often referred to as “motivational judgment internalism,” see Rosati (2016). For a prominent defense of the view, see Smith (1994). Of course, there are related but distinct internalist views in the literature, but the broad outlines will suffice for my purposes here.

this case, by looking at the limitations of what human beings are capable of, including psychological limitations, we can learn about the extent of our moral obligations.

This is all to make the point that what I have argued here is tied to well-established tendencies in moral and political philosophy and a variety of views and positions in these areas. In terms of the notion of fruitfulness that I have described, there may also be other ways in which moral and political concepts can help us to solve practical problems, such as fruitfulness in helping us to classify various actions, and in allowing us to justify various assignments of duties and responsibilities. Much more can be said on these topics, but the point to note here is that I am interested in particular types of practical fruitfulness, and do not deny that there may be other types of fruitfulness that can go into the evaluation of moral and political concepts.

Further, as I have suggested above, no claim regarding a concept's fruitfulness is likely to be established definitively by the results of only one study. Nor does any one empirical field have a claim to fundamentality or pre-eminence in providing evidence regarding the fruitfulness of a given normative concept in solving a particular practical problem. Rather, evidence must be taken from both the armchair and the lab and balanced in reflective equilibrium to determine which normative concepts and their conceptions can best solve practical problems. It is in this way that the unspoken agreement that I mentioned at the beginning of Section 1, between experimental philosophers and moral and political philosophers who are receptive to empirical research but do not themselves conduct any, can be vindicated.

3. Fruitfulness, Theoretical Criteria, and Practical Problems

I mentioned earlier that fruitfulness is a criterion used to evaluate normative concepts in addition to their being coherent and well motivated from a purely theoretical perspective. In the case of arguments, we are generally only interested in the fruitfulness of ones that we take to be valid and sound. The exception is when an argument contains a premise that makes a claim about the fruitfulness of a concept, in which case the argument would be unsound if we found that claim to be false. But fruitfulness is not meant to shore up concepts or arguments that are bad by purely theoretical standards, ignoring the fact that they are incoherent, invalid, or unsound, as the case may be. The point is to see which concepts and arguments, that at least some reasonable interlocutors take to be good ones, do the best job of helping us to solve practical problems, and of course also to evaluate claims to fruitfulness that might be used in philosophical arguments.

Further, this should not be taken to mean that just any practical problem that a philosophical concept helps us to solve would be relevant to the evaluation of that concept. If the internalization of a philosophical concept unrelated to teenage smoking somehow led to lower teen smoking rates in a population, that would not be a practical problem whose resolution would plausibly bear a relation to the concept's evaluation from the standpoint of philosophy (though perhaps it would from the standpoint of public health). So *relevance* is a condition relating philosophical concepts to practical problems on the basis of subject matter that must be met for the evaluation of a concept in terms of its effects on a practical problem to be philosophically important. My own view is that while some obviously irrelevant practical problems could be impacted by normative concepts, we will generally be able to distinguish

these problems from practical problems that are in the ballpark of the concept we are considering. Within this range of problems, we should be fairly liberal in examining the effects that normative concepts can have.

Additionally, the *way* in which a moral or political concept solves a practical problem is important, and not just any way will do. It would be a mistake to take the view that normative concepts can be evaluated in terms of fruitfulness to mean that these concepts are mere coordination mechanisms, to be evaluated solely in terms of whether they bring about stability, motivate people to perform the right kinds of actions, or meet some other practical desiderata. The way that they solve practical problems must be what I will call “value-consistent.” This means first that the way that the concept solves the relevant practical problem must be consistent with other plausibly moral values, as opposed to non-moral values, such as control, domination, and so on. Rawls refers to this criterion, in the case of a theory of justice, as that of achieving “stability for the right reasons” (Rawls 1993). And second, the normative concept must not violate its own commitments in solving a practical problem. To return to the same well-known example, a liberal theory of justice cannot violate the norms of liberalism, which include the commitment that a theory of justice must be neutral among a range of reasonable worldviews or “comprehensive doctrines,” in solving the problem of distributive justice in a democratic society. This second criterion doesn’t depend on liberalism itself being the right political doctrine, but rather on coherence with liberalism’s own commitments. Each of these conditions must be met for the way a normative concept solves a practical problem to be normatively acceptable and should be included in our best conception of what makes normative concepts fruitful.

4. Carnapian Theory Construction and Fruitfulness

Joshua Shepherd and James Justus have also recently defended the positive program in experimental philosophy (Shepherd and Justus 2015), in their case drawing heavily on Rudolf Carnap's views on theory construction (Carnap 1950, 1955). Their defense focuses on Carnap's description of the first of three stages of research that should be conducted in building a philosophical theory, which they call "explication preparation," where the other two stages are explication and fruitfulness. My own view borrows the term it focuses on from Carnap's third stage, although my understanding of fruitfulness differs greatly from Carnap's own. It will be worth pausing to reflect on this difference, and also how Shepherd and Justus' project differs from mine.

Shepherd and Justus argue that experimental philosophy can play an important role in explication preparation, and thereby play a positive role in developing philosophical theories. For Carnap, good philosophical theorizing involves the "explication" of concepts, which consists of identifying, evaluating, and typically revising the contents of ordinary concepts that we possess and use. Philosophers precisify ordinary concepts in order that they can better serve purposes, primarily the purposes of scientific inquiry for Carnap. Explication preparation involves clarifying the content of an ordinary concept in advance of explication. Once these raw materials have been clarified, explication can proceed in precisifying the concept that is to be used in scientific investigation. Then, that explication or precisification can be evaluated in terms of whether it enhances "fruitfulness." Carnap's core example of the assessment of an explication in terms of fruitfulness is the move from the ordinary concept 'fish' to the scientific concept 'piscis' (Carnap 1950, pp. 5–6). The concept 'fish' is vague and includes many animals

that are not cold-blooded or do not have gills throughout their lives. The concept ‘piscis’ instead denotes only aquatic animals that are cold-blooded and that do have gills throughout life. This precisified concept is more fruitful in virtue of allowing a larger number of true general statements to be formulated, thereby helping us to advance our understanding of underwater life.

Shepherd and Justus focus on explication preparation as the point of contact between experimental philosophy and traditional philosophical methods. Empirical research can reveal ways in which the extensions or intensions of concepts are vague, show that there are multiple concepts operating in our thinking about a certain notion or topic that we took to primarily involve only one concept, discover biases influencing intuitive judgments and their sources, discover other unpredictable, non-bias-related influences on judgments about concepts, and highlight the central features of a concept and any dependence relationships that concept stands in with others (Shepherd and Justus 2015, pp. 390-1). As they acknowledge, empirical research will not fully determine how any particular explication should proceed. Explication involves choices that are partly guided by the theoretical aims that we have, such as choices about which aspects of a concept must be held onto and which can be abandoned. But nonetheless, there is a positive contribution that empirical research can make to philosophy, according to Shepherd and Justus, in aiding explication preparation.

Shepherd and Justus have provided an important account of one role that experimental research can play in philosophical theory construction. However, two important points are worth noting that should give us pause regarding their account. First, it relies fairly heavily on the Carnapian view of theory construction, and insofar as that view is outmoded or controversial, their account will rest on somewhat shaky footing. Second, and crucially, their view keeps experimental philosophy firmly in a handmaiden position with respect to philosophical research,

securing the preconditions for good philosophical theory construction without participating in it. Part of the challenge for experimental philosophy is to explain why it is *philosophy*, not just helpful for preventing or correcting some mistakes that we might make in the early stages of developing our philosophical theories.

Rather than *preparing* us to do philosophy, my view has the implication that experimental philosophy is part of the philosophical enterprise. Instead of Carnap's notion of fruitfulness, which involves yielding a larger number of true general statements, I have presented a number of ways in which philosophical concepts can be fruitful in helping us to solve practical problems. These types of fruitfulness, I have noted, are tied to practical roles that moral and political concepts have been expected by philosophers to be evaluated in relation to. Insofar as one of these philosophical concepts is fruitful, that is a strike in its favor *qua* philosophical concept.

This is not to say that I don't think that experimental philosophy can play the role that Justus and Shepherd describe. Nor am I taking a stand here on the role empirical research might have in explication, or its analogues in other views of theory construction. My view centers on a plausible notion of fruitfulness for moral and political concepts that fits well with the practical roles that many philosophers have already thought that these concepts ought to play. Because these concepts should be evaluated, in part, by criteria of fruitfulness, and empirical research will be important in determining whether or not concepts fulfill these criteria, empirical research is part of the evaluation of these concepts, and hence part of the philosophical enterprise.

5. Haslanger's Ameliorative Inquiry and Resilience

I should also discuss Sally Haslanger's influential view about concepts and the different methods that we might take up in engaging with them (Haslanger 2000, 2012) and how her view relates to what I have said here. Haslanger argues that some concepts might play roles that need to be filled, for instance in a feminist antiracist theory, and that this might be a way of vindicating these concepts. What she calls an "ameliorative inquiry" into questions concerning the nature of gender and race examines the point of having these concepts and what cognitive or practical problems they might be retooled to help us solve. Our concepts serve cognitive and practical purposes, and these purposes can be advanced or held back depending on how these concepts are characterized. Ameliorative inquiry examines not only the practical roles that concepts in fact play, but also what roles they could play for us in addressing injustices, such as gender inequality and racism. Her approach is meant to go beyond the method of conceptual analysis characteristic of ordinary language philosophy, which might examine concepts like gender and race by attempting to determine the necessary and sufficient conditions of their applications as we currently use them. It is also distinct from a descriptive project that looks at the extension of concepts and, perhaps using empirical means, seeks to determine whether or not these concepts pick out social kinds. Rather, ameliorative inquiry is supposed to allow us to revise our conceptual practices and the extensions of our concepts to better serve the moral or political goals that they can help us to advance.

Some of the points that I have made in this paper can be seen as emphasizing the role that empirical research should play in ameliorative inquiry. As mentioned above, Haslanger's goal of determining which concepts of gender and race "will be effective tools in the fight against

injustice” (2000, p. 36) fits well with my characterization of the resilience-based fruitfulness of concepts. Ameliorative inquiry is, in a sense, a project of determining whether important concepts like gender and race can be improved upon in terms of one type²¹ of fruitfulness.

Haslanger is surely right that seemingly non-moral concepts can play political roles and help to structure our interactions. Hence, I don’t want my focus on normative concepts to be taken as a point of disagreement. Rather, my area of greatest familiarity is with normative concepts. So it is worth distinguishing between how non-normative concepts might be fruitful in a normative way—realizing various moral values or helping to produce political outcomes—without themselves being normative concepts, and the fruitfulness of normative concepts. I take it that Haslanger is more focused on the former, and I am more focused on the latter. Still, it seems to me that much of what I have to say about how normative concepts can be fruitful may also apply to some non-normative concepts. Filling in some of these details is part of the work of my paper that may also be applied as a further defense of Haslanger’s notion of ameliorative inquiry, and how empirical research can aid such inquiry. Of course, Knobe’s research has shown us that it may be difficult to distinguish between normative and non-normative concepts if our criterion is whether normative considerations go into their application.²² Subject matter seems more appropriate as a criterion, where normative concepts focus on normative subjects, such as morality and aesthetics, and non-normative concepts focus on subjects that are typically descriptive. On my view, moral and political concepts should be assessed, in part, in terms of

²¹ It may be the case that the only concepts of gender and race that are fruitful in the resilience sense, of course, will also have to be fruitful in other respects – motivating various kinds of feminist, antiracist actions, preventing sexist and racist actions, and so on. As with other concepts, an exploration of the resilience-based fruitfulness of gender and race may lead to questions concerning the other types of fruitfulness that I have described.

²² For a helpful overview, see Knobe (2010).

their fruitfulness in helping us to solve practical problems. This view is fully compatible with the view that other concepts may also play important practical roles in political life, and therefore also be assessed in terms of their fruitfulness.

6. The Role of Philosophers in Empirical Research

But why should philosophers be the ones conducting experiments? Why isn't everything that I've said best taken as suggesting that philosophers should engage with the work of, and perhaps collaborate with, psychologists? Depending on what is meant by "collaboration," I will disagree or agree with the statement implied by the second question. If collaboration means discussing the experiments and their design, talking through the relevant philosophical issues, and working to interpret the data and its philosophical significance, there is nothing wrong with the statement. This kind of "active" collaboration, to give it a name, is fine and to be encouraged even. The problem is "passive" or inactive collaboration, where the philosopher simply draws on existing empirical research, typically only the research that seems to fit with what they want to say, and leaves the rest of the empirical issues to the psychologists. Indirectly, this problem of "cherry picking" will be more likely to be avoided if philosophers are actively collaborating with psychologists and working on empirical research. But the further answer to the first question of why philosophers shouldn't just let others conduct the experiments, in the sense of being involved in their design, administration, and interpretation, is that the right kinds of empirical work won't get done on this approach. Philosophers are trained to draw distinctions that psychologists are generally not sensitive to. Our disciplinary foci are different, and the ornateness of philosophical questions makes it the case that philosophical training is generally

required to address them. Training is required to appreciate the importance of the nitty gritty questions that are within a philosophical debate, and how to design experiments to test questions at a level of specificity that is germane to those parts of a debate. In general, philosophical work isn't on the biggest questions, but rather takes place within a set of assumptions and conversations that presuppose particular ways of approaching those questions. With the questions raised in moral and political philosophy, it is important that people who have the training to see just what concepts are being evaluated and how their fruitfulness is relevant to the philosophical questions under consideration be involved in designing, conducting, and interpreting the results of empirical research. On my view, evaluating the fruitfulness of normative concepts is part of the enterprise of doing moral and political philosophy. But conducting empirical research to determine the fruitfulness of these concepts and assessing the relevance of this research to philosophical debates are activities that involve attention to distinctions and subtleties that generally requires philosophical training. An important upshot of my view is that moral and political philosophers must be actively involved in conducting empirical research that will help us to ascertain whether particular normative concepts are fruitful ones.

7. Conclusion

Moral and political concepts can be evaluated, in part, in terms of their fruitfulness. What kinds of actions they motivate people to perform and prevent them from performing, whether they help us in fighting back against injustices and other problematic social phenomena, whether they can be subjects of consensus among people of good will, and whether they provide

us with useful guidance are factors that determine whether these concepts can help us to solve practical problems. The view that I have presented here shows that experimental philosophy can be part of normative philosophical research, because assessing the fruitfulness of moral and political concepts will require empirical research. Rather than merely clearing the way for traditional methods to be employed, empirical research is continuous with traditional philosophical methodology. This view is also compatible with a plausible explanation of why philosophers, in particular, should be involved in empirical research. Of course, no interesting view appeals to everyone. Some philosophers are opposed to letting empirical research into the set of approaches that shed light on philosophical questions. Yet many other philosophers, as I have shown, should be receptive to the view that I have offered. Indeed, as long as one regards “what practical purposes should our moral and political concepts serve?” as a sensible question to ask, they are the sort of theorist who can endorse what I have said about the value of empirical research in addressing issues in moral and political philosophy.²³

²³ For early conversations on the topics addressed in this paper, I am grateful to Serene Khader, Josh Knobe, Tori McGeer, and Michael Strevens. For helpful discussions of earlier versions of the paper and the points raised in it, I would like to thank Christian Barry, Geoff Brennan, Ramon Das, Ben Fraser, Ana Gantman, Simon Keller, Eric Mandelbaum, Philip Pettit, Luke Russell, Nic Southwood, Katie Steele, Kim Sterelny, Justin Sytsma, John Thrasher, Joe Ulatowski, Dan Weijers, and audiences at the First Annual Australasian Experimental Philosophy Workshop at Victoria University of Wellington, the Fifth Australasian Workshop in Moral Philosophy at the Kioloa Coastal Campus of the Australian National University, the Experimental Methods in Social and Political Philosophy Workshop at ANU, and the ANU School of Philosophy Seminar. For helpful written comments, I am grateful to Justin Bruner, Toby Handfield, Frank Jackson, Serene Khader, Josh Knobe, Matt Kopec, Shaun Nichols, Peter Singer, Katja Vogt, and two anonymous reviewers for the journal.

References

- Alexander, J., Mallon, R., & Weinberg, J. M. (2014). Accentuate the Negative. In J. Knobe & S. Nichols (Eds.), *Experimental Philosophy* (Vol. 2, pp. 31–50). Oxford: Oxford University Press.
- Alexander, J., & Weinberg, J. M. (2007). Analytic Epistemology and Experimental Philosophy. *Philosophy Compass*, 2(1), 56–80.
- Amodio, D.M., Devine, P.G., & Harman, E. (2008). Individual Differences in the Regulation of Intergroup Bias: The Role of Conflict Monitoring and Neural Signals for Control. *Journal of Personality and Social Psychology*, 94, 60-74.
- Anderson, E. (2010). *The Imperative of Integration*. Princeton, NJ: Princeton University Press.
- Aristotle. (1999). *Nicomachean Ethics*. (T. Irwin, Ed.) (2nd ed.). Indianapolis, Ind: Hackett Pub. Co.
- Barrett, L.F., Gollwitzer, P.M., Lee-Chai, A., Barndollar, K., & Trötschel, R. (2001). Individual Differences in Working Memory Capacity and Dual-Process Theories of Mind. *Psychological Bulletin*, 130, 553-575.
- Bateson, M., Nettle, D., & Roberts, G. (2006). Cues of Being Watched Enhance Cooperation in a Real-World Setting. *Biology Letters*, 2(3), 412–414.
- Batson, C. D. (2008). Moral Masquerades: Experimental Exploration of the Nature of Moral Motivation. *Phenomenology and the Cognitive Sciences*, 7(1), 51–66.
- Bonilla-Silva, E. (2003). *Racism without Racists: Color-Blind Racism and the Persistence of Racial Inequality in America*. Lanham, MA: Rowman & Littlefield.
- Boxill, B. R. (1992). *Blacks and Social Justice* (Rev. ed.). Lanham, Md: Rowman & Littlefield.
- Bruner, J. (2018). Decisions Behind the Veil: An Experimental Approach. In T. Lombrozo, J. Knobe, & S. Nichols (Eds.), *Oxford Studies in Experimental Philosophy* (Vol. 2). Oxford: Oxford University Press.
- Bruner, J. and Lindauer, M. (forthcoming) The Varieties of Impartiality, or, Would an Egalitarian Endorse the Veil? *Philosophical Studies*.
- Bryan, C.J., Walton, G.M., Rogers, T., & Dweck, C.S. (2011). Motivating Voter Turnout by Invoking the Self. *Proceedings of the National Academy of Sciences*, 108(31), 12653-12656.

- Bryan, C.J., Adams, G.S., & Monin, B. (2012). When Cheating Would Make You a Cheater: Implicating the Self Prevents Unethical Behavior. *Journal of Experimental Psychology: General*, 142(4), 1001-1005.
- Buckland, L., Lindauer, M., Rodríguez-Arías, D., and Véliz, C. Testing the Motivational Strength of Positive and Negative Duty Arguments Regarding Global Poverty. Unpublished manuscript.
- Cappelen, H. (2012). *Philosophy Without Intuitions*. Oxford: Oxford University Press.
- Carnap, R. (1950). *Logical Foundations of Probability* (Second Edition.). Chicago: University of Chicago Press.
- Carnap, R. (1955). Meaning and Synonymy in Natural Languages. *Philosophical Studies*, 6(3), 33–47.
- Cohen, G. A. (2008). *Rescuing Justice and Equality*. Cambridge, Mass: Harvard University Press.
- Correll, J., Park, B., Judd, C. M., & Wittenbrink, B. (2002). The Police Officer's Dilemma: Using Ethnicity to Disambiguate Potentially Threatening Individuals. *Journal of Personality and Social Psychology*, 83(6), 1314–1329.
- Cushman, F. (2013). Action, Outcome, and Value: A Dual-System Framework for Morality. *Personality and Social Psychology Review*, 17(3), 273-292.
- Enoch, D. (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Ernest-Jones, M., Nettle, D., & Bateson, M. (2011). Effects of Eye Images on Everyday Cooperative Behavior: A Field Experiment. *Evolution and Human Behavior*, 32(3), 172–178.
- Essien, I., Stelter, M., Kalbe, F., Koehler, A., Mangels, J., & Meli, S. (2017). The Shooter Bias: Replicating the Classic Effect and Introducing a Novel Paradigm. *Journal of Experimental Social Psychology*, 70, 41–47.
- Estlund, D. (2014). Utopophobia. *Philosophy & Public Affairs*, 42(2), 113–134.
- Feinberg, M., & Willer, R. (2015). From Gulf to Bridge: When Do Moral Arguments Facilitate Political Influence? *Personality and Social Psychology Bulletin*, 1-17.
- Fisher, J. C. (2015). Pragmatic Experimental Philosophy. *Philosophical Psychology*, 28(3), 412–433.

- Frohlich, N., Oppenheimer, J. A., & Eavey, C. L. (1987). Laboratory Results on Rawls's Distributive Justice. *British Journal of Political Science*, 17(01), 1.
- Garner, R. (2007). Abolishing Morality. *Ethical Theory and Moral Practice*, 10(5), 499–513.
- Gaus, G. (2010). *The Order of Public Reason: A Theory of Freedom and Morality in a Diverse and Bounded World*. Cambridge, UK: Cambridge University Press.
- Graham, J., Nosek, B.A., Haidt, J., Iyer, R., Spassena, K., & Ditto, P.H. (2011). Mapping the Moral Domain. *Journal of Personality and Social Psychology*, 101(2), 366–385.
- Greene, J. D. (2007). The Secret Joke of Kant's Soul. In W. Sinnott-Armstrong (Ed.), *Moral Psychology, Vol. 3: The Neuroscience of Morality: Emotion, Disease, and Development*. Cambridge, MA: MIT Press.
- Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science (New York, N.Y.)*, 293(5537), 2105–2108.
- Haidt, J. (2001). The Emotional Dog and its Rational Tail: A Social Intuitionist Approach to Moral Judgment. *Psychological Review*, 108(4), 814–834.
- Haidt, J. (2012). *The Righteous Mind*, New York: Pantheon Books.
- Haslanger, S. (2000). Gender and Race: (What) Are They? (What) Do We Want Them To Be? *Noûs*, 34(1), 31–55.
- Haslanger, S. (2012). *Resisting Reality: Social Construction and Social Critique*. New York: Oxford University Press.
- Hassoun, N., Malikov, E., & Lubchenco, N. (2016). How People Think About Distributing Aid. *Philosophical Psychology*, 29(7), 1029–1044.
- Jaggar, A. M. (2001). Is Globalization Good for Women? *Comparative Literature*, 53(4), 298.
- Jaggar, A. M. (2005). "Saving Amina": Global Justice for Women and Intercultural Dialogue. *Ethics & International Affairs*, 19(03), 55–75.
- Kant, I. (2012). *Groundwork of the Metaphysics of Morals*. (M. J. Gregor & J. Timmermann, Eds.) (Revised Edition). Cambridge: Cambridge University Press.
- Kennett, J., & Fine, C. (2009). Will the Real Moral Judgment Please Stand Up?: The Implications of Social Intuitionist Models of Cognition for Meta-ethics and Moral Psychology. *Ethical Theory and Moral Practice*, 12(1), 77–96.
- Khader, S. J. (2011). *Adaptive Preferences and Women's Empowerment*. New York: Oxford University Press.

- Kitcher, P. (2014). *The Ethical Project*. Cambridge, Mass.: Harvard Univ. Press.
- Knobe, J. (2010). Person as Scientist, Person as Moralist. *Behavioral and Brain Sciences*, 33(04), 315–329.
- Knobe, J. (2016). Experimental Philosophy Is Cognitive Science. In J. Sytsma & W. Buckwalter (Eds.), *A Companion to Experimental Philosophy* (pp. 37–52). Chichester, UK: John Wiley & Sons, Ltd.
- Korsgaard, C. M. (2003). Realism and Constructivism in Twentieth-Century Moral Philosophy. *Journal of Philosophical Research*, 28(Supplement), 99–122.
- Lawford-Smith, H. (2012). The Motivation Question: Arguments from Justice and from Humanity. *British Journal of Political Science*, 42(03), 661–678.
- Lerner, J. S., & Tetlock, P. E. (1999). Accounting for the Effects of Accountability. *Psychological Bulletin*, 125(2), 255–275.
- Levy Paluck, E., Shepherd, H. & Aronow, P.M. (2016). Changing Climates on Conflict: A Social Network Experiment in 56 Schools. *Proceedings of the National Academy of Sciences*, 113(3), 566-571.
- Lichtenberg, J. (2010). Negative Duties, Positive Duties, and the “New Harms.” *Ethics*, 120(3), 557–578.
- Lichtenberg, J. (2014). *Distant Strangers: Ethics, Psychology, and Global Poverty*.
- Lissowski, G., Tyszka, T., & Okrasa, W. (1991). Principles of Distributive Justice: Experiments in Poland and America. *Journal of Conflict Resolution*, 35(1), 98–119.
- Marx, D. M., & Roman, J. S. (2002). Female Role Models: Protecting Women’s Math Test Performance. *Personality and Social Psychology Bulletin*, 28(9), 1183–1193.
- McIntyre, R. B., Lord, C. G., Gresky, D. M., Ten Eyck, L. L., Frye, G. D. J., & Bond Jr., C. F. (2005). A Social Impact Trend in the Effects of Role Models on Alleviating Women’s Mathematics Stereotype Threat. *Current Research in Social Psychology*, 10, 116–136.
- McIntyre, R. B., Paulson, R. M., & Lord, C. G. (2003). Alleviating Women’s Mathematics Stereotype Threat Through Salience of Group Achievements. *Journal of Experimental Social Psychology*, 39(1), 83–90.
- Mikhail, J. M. (2007). Universal Moral Grammar: Theory, Evidence and the Future. *Trends in Cognitive Sciences*, 11(4), 143–152.

- Mikhail, J. M. (2011). *Elements of Moral Cognition: Rawls' Linguistic Analogy and the Cognitive Science of Moral and Legal Judgment*. Cambridge ; New York: Cambridge University Press.
- Mills, C. W. (2005). "Ideal Theory" as Ideology. *Hypatia*, 20(3), 165–184.
- Mills, C. W. (2011). *The Racial Contract* (Nachdr.). Ithaca, NY: Cornell Univ. Press.
- Moehler, M. (2018). *Minimal Morality: A Multilevel Social Contract Theory*. Oxford University Press.
- Muldoon, R., Lisciandra, C., Colyvan, M., Martini, C., Sillari, G., & Sprenger, J. (2014). Disagreement behind the veil. *Philosophical Studies*, 170(3), 377-394.
- Nadelhoffer, T., & Nahmias, E. (2007). The Past and Future of Experimental Philosophy. *Philosophical Explorations*, 10(2), 123–149.
- Parfit, D. (1992). *Reasons and Persons* (Repr.). Oxford: Clarendon.
- Payne, B. K. (2005). Conceptualizing Control in Social Cognition: How Executive Functioning Modulates the Expression of Automatic Stereotyping. *Journal of Personality and Social Psychology*, 89(4), 488–503.
- Pettit, P. (1999). *Republicanism: A Theory of Freedom and Government*. Oxford: Oxford University Press.
- Pettit, P. (2014). *Just Freedom: A Moral Compass for a Complex World* (First Edition). New York: W.W. Norton & Company.
- Pinder, M. (forthcoming). Does Experimental Philosophy Have a Role to Play in Carnapian Explication?: X-Phi and Explication. *Ratio*.
- Pogge, T. W. (2008). *World Poverty and Human Rights: Cosmopolitan Responsibilities and Reforms* (Second Edition). Cambridge: Polity Press.
- Rawls, J. (1971). *A Theory of Justice* (rev. 1999.). Cambridge, Mass: Belknap Press of Harvard University Press.
- Rawls, J. (1980). Kantian Constructivism in Moral Theory. *The Journal of Philosophy*, 77(9), 515.
- Rawls, J. (1993). *Political Liberalism*. New York: Columbia University Press.
- Rawls, J. (2001). *Justice as Fairness: A Restatement*. Cambridge, Mass: Harvard University Press.

- Rosati, C. S. (2016, Winter Edition). Moral Motivation. In (E. N. Zalta, Ed.) *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/win2016/entries/moral-motivation>
- Shepherd, J., & Justus, J. (2015). X-Phi and Carnapian Explication. *Erkenntnis*, 80(2), 381–402.
- Sidgwick, H. (1907). *The Methods of Ethics* (Seventh Edition). London: Macmillan.
- Singer, P. (1972). Famine, Affluence, and Morality. *Philosophy and Public Affairs*, 1(3), 229–243.
- Singer, P. (2005). Ethics and Intuitions. *The Journal of Ethics*, 9(3–4), 331–352.
- Skitka, L.J., Bauman, C.W., & Sargis, E.G. (2005). Moral Conviction: Another Contributor to Attitude Strength or Something More? *Journal of Personality and Social Psychology*, 88(6), 895-917.
- Skitka, L.J. (2010). The Psychology of Moral Conviction. *Social and Personality Psychology Compass*, 4(4), 267-281.
- Skyrms, B. (1996). *Evolution of the Social Contract*. Cambridge University Press.
- Skyrms, B. (2016). Evolution, Norms and the Social Contract. *Arizona State Law Journal*, 48, 1087-1100.
- Slovic, P. (2007). “If I Look at the Mass I Will Never Act”: Psychic Numbing and Genocide. *Judgment and Decision Making*, 2(2), 79–95.
- Small, D. A., Loewenstein, G., & Slovic, P. (2007). Sympathy and Callousness: The Impact of Deliberative Thought on Donations to Identifiable and Statistical Victims. *Organizational Behavior and Human Decision Processes*, 102, 143-153.
- Smith, M. (1994). *The Moral Problem*. Oxford, UK ; Cambridge, Mass., USA: Blackwell.
- Southwood, N. (2016). Does “Ought” Imply “Feasible”? *Philosophy & Public Affairs*, 44(1), 7–45.
- Street, S. (2008). Constructivism about Reasons. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. 3, pp. 207–245). Oxford: Oxford University Press.
- Street, S. (2016). Constructivism in Ethics and the Problem of Attachment and Loss. *Aristotelian Society Supplementary Volume*, 90(1), 161–189.
- Sugden, R. (1990). Contractarianism and Norms. *Ethics*, 100(4), 768-786.
- Sytsma, J., & Livengood, J. (2016). *The Theory and Practice of Experimental Philosophy*. Peterborough, Ontario: Broadview Press.

- Thrasher, J. (2012). Uniqueness and symmetry in bargaining theories of justice. *Philosophical Studies*, 167(3), 683-699.
- Valentini, L. (2012). Ideal vs. Non-ideal Theory: A Conceptual Map: Ideal vs Non-ideal Theory. *Philosophy Compass*, 7(9), 654–664.
- Vohs, K. & Schooler, J.W. (2008) The Value of Believing in Free Will: Encouraging a Belief in Determinism Increases Cheating. *Psychological Science*, 19(1), 49-54.
- Walker, M.U. (2007). *Moral Understandings: A Feminist Study in Ethics* (Second Edition). Oxford: Oxford University Press.
- Weinberg, J. M. (2007). How to Challenge Intuitions Empirically Without Risking Skepticism. *Midwest Studies in Philosophy*, 31(1), 318–343.
- Wicklund, R.A. (1979). The Influence of Self-Awareness on Human Behavior: The Person Who Becomes Self-Aware is More Likely to Act Consistently, Be Faithful to Societal Norms, and Give Accurate Reports About Himself. *American Scientist*, 67(2), 187-193.
- Williamson, T. (2016). Philosophical Criticisms of Experimental Philosophy. In J. Sytsma & W. Buckwalter (Eds.), *A Companion to Experimental Philosophy* (pp. 22–36). Chichester, UK: John Wiley & Sons, Ltd.
- X-Phi's New Take on Old Problems. (2010, August 19). *The New York Times*.
<https://www.nytimes.com/roomfordebate/2010/8/19/x-phis-new-take-on-old-problems>