

The Democratic Limits of Political Experiments

Eric Beerbohm (Beerbohm@fas.harvard.edu)

Ryan Davis (rwdavis@byu.edu)

Adam Kern (abkern@princeton.edu)

Experiments *on* campaigns go back almost a hundred years, when political scientists conducted the first randomized field experiment testing voter support of Calvin Coolidge. Experiments *by* campaigns are an entirely new feature of democratic politics, as close to a “new moral fact” as we have seen in the electoral setting.¹ Yet neither intervention — by political science or campaigns — has faced the scrutiny of democratic theorists. In this paper we begin to remedy this inattention. Is the victory lab a boon or loss to our ideal of democracy?

Field experiments in democratic politics are freighted with normative significance. They do not only allow researchers and campaign staffs to collect data. They also influence the actions and relationships of various political officers—including voters. Any intentional influencing of morally relevant actions and relationships will carry moral questions in tow, and political experiments are no exception. Do citizens and/or elite political actors possess moral complaints against certain field experiments? If there is a conflict between democratic values and intellectual ones, how should these concerns be traded off? Are there moral limits on how experimenters might permissibly influence the rational deliberation of political actors?

¹ For the notion of a “new moral fact,” see Herman 1993.

These questions have not been the subject of sustained philosophical scrutiny.² To be sure, there are extensive philosophical literatures on experimental ethics in other fields. The literature on medical ethics, for example, focuses on issues of consent and deception within clinical trials. Insights from these fields will certainly help to develop a full view of the ethics of political field experiments. But we think that the political context—and the democratic context, in particular—raises distinctive moral concerns, which the ethics literatures from other fields do not capture. Meanwhile, several practitioners of field experiments have begun to reflect, quite sensitively, on the ethics of their work. But these efforts have not been deeply grounded in the rich literatures of political philosophy and normative ethics.

In this paper, we hope to fill this gap: we hope to articulate some distinctively political concerns about field experiments, in a relatively systematic fashion. Our goal is not to *vindicate* these concerns—let alone charge anyone with wrongdoing. It is rather to clarify just what these concerns are and identify what further reflection is needed to adjudicate them.

We will argue for three claims. First, campaign experiments could be objectionable even if they do not coerce or deceive anyone. Second, campaign experiments could be objectionable even if they are not *decisive*—even if they do not swing the results of elections or other political decisions. Third, if campaign experiments *are* objectionable in these ways, one reason for this is that they undermine citizens' equal standing.

² This is importantly not to say that these questions have avoided ethically motivated critical attention. Our essay will follow a variety of practitioners who have already begun the process of self-conscious reflection on the ethics of experiments in political science. See Gerber 2011; Levitt and List 2009; Teele 2014; Barrett and Carter 2014; and the contributions to Desposato 2016. There has also been extensive work by political philosophers about the values associated with democratic communities, including Beitz 1989; Gutman and Thompson 1996; Christiano 2008; Kolodny 2014a; Viehoff 2014; and Rose 2016. Although political experiments implicate democratic values, these two literatures have not yet been brought into conversation with each other.

Here we limit our argument to field experiments within electoral settings. We do this for two reasons. First, given the scope of experiments in politics, including experiments on diverse sets of officials, this limiting condition will help streamline our argument and the relevant experiments (though even in this universe, we have found hundreds of experiments). Second, elections are arguably the purest case in politics, representing a relatively closed system with a clear endpoint; in this sandbox, the effects of experimental treatments are comparatively easy to measure. That being said, we don't think that the upshot of our argument will be limited to this setting. The moral concerns that we articulate, if they are justified, can be raised throughout democratic politics.

This paper proceeds in the following sections. After performing preliminaries, we identify some recent experiments that are morally salient. To say that they are morally salient is not to register any judgments about their permissibility or impermissibility, justice or injustice, or the like. It is merely to observe that they call out for further moral consideration. Second, we develop a theory of what might be democratically objectionable about these experiments. Third, we consider the implications of this theory. Most critically, our theory does not take us to extremes: it does not commit us to claiming that all experiments are forbidden or that all are permitted. Finally, we consider some ways that our account might be extended outside of the electoral setting.

1. Preliminaries: Manipulation(s)

We begin by clarifying both our methodology and a central term: "manipulation." Both our methodology and what we shall mean by "manipulation" follow from the aim that we pursue in this paper. Our aim is not to charge anyone conducting field experiments with

moral wrongdoing. Rather, it is to clarify the logical space of potential moral concerns, allowing for more systematic consideration of these features. We hope to chart the moral terrain, not to climb any of its summits.

Thus, our methodology is to use the standard item in the philosopher’s toolkit—reflective equilibrium—though with a somewhat more expansive twist. Reflective equilibrium is the generally Socratic method, formalized by John Rawls, of collecting one’s judgments about cases and theories, forming a set of principles to account for them, and then attempting to reconcile those principles with one’s judgments, until one comes to “reflective equilibrium” (Rawls [1951] 1999b; [1971] 1999a; [1974] 1999b; [1980] 1999b; [1993] 1996). Since we are prioritizing the articulation of various moral concerns, rather than their conclusive evaluation, we will modify this method in two respects. First, we will give somewhat more initial credence to initial *negative* judgments about cases—that is, judgments that the actions considered are wrong, or unjust, or something of that sort.³ Second, we will break off the process of reflective equilibrium before coming to conclusive judgments about whether the actions we consider *are* wrong, unjust, and so on. As we will show, forming such conclusive judgments about field experiments would require resolving some deep issues which are unresolved in the philosophical literature. Unfortunately, given space constraints, we simply cannot both taxonomize possible democratic concerns about

³ Our method thus reverses the customary weight accorded to Type I versus Type II errors in political science (Cf. Zimmerman 2016, p. 191). However, this reversal follows from differing attitudes towards risk, which make sense in light of our aims and the aims of political scientists. Since we are interested in canvassing potential moral objections, we are especially concerned about *overlooking* potentially valid ones. For social scientists, meanwhile, the danger of mistakenly *accepting* an hypothesis is of greater concern.

field experiments and then conclusively evaluate them. Here we perform the taxonomy—the logically prior task—and defer the conclusive evaluation for future work.⁴

We will use “manipulation” as a residual term to describe all intentional influencing of another’s behavior which is both (a) morally objectionable and (b) neither coercive nor deceptive. We adopt this stipulative definition for two main reasons. First, since manipulation is a fiercely contested concept in the philosophical literature, we cannot embrace any more specific definition of it uncontroversially (Baron 2003; Barnhill 2014; Buss 2005; Cave 2007; Conly 2004; Greenspan 2003; Hill 1991; Manne 2014). Second, again, this definition accords with our underlying aim (Cf. Godfrey-Smith 2012). Since we are more concerned about neglecting genuine moral concerns than giving credence to illusory ones, we prefer to err on the side of an expansive definition of manipulation rather than a restrictive one.

It is important to note that the philosopher’s concept of manipulation differs from the social scientist’s. “Manipulation”—understood as an intervention by the researcher that creates a difference in circumstances between treatment and control groups—is, of course, characteristic of experimental work in general. When Teele (2014, 118) suggests that “the term manipulation seems rather anodyne,” she has in mind the social scientist’s concept, and not the philosopher’s. We should not be read as suggesting at any point that manipulation, in the social scientist’s sense, is generally morally objectionable.

⁴ For an example of such a two-part treatment of another subject, campaign finance, see Pevnick 2016a; Pevnick 2016b.

Still, we recognize that some will resist our definition of manipulation.⁵ While we acknowledge the inevitability of some conceptual controversy, we believe that our working version of the concept genuinely helps to clarify our subject matter, helps us to achieve our theoretical aims, and remains acceptable to as wide an audience as possible. Readers might feel dissatisfied that our conception of manipulation does not do more to pick out cases that are morally concerning. Instead, by providing a moralized definition (that is, by including moral concern as part of the definition, itself), our account cannot—by itself—explain what is morally objectionable about manipulation. This is a limitation. But the aspiration to give unified, non-moralized accounts of similar phenomena have largely met with philosophical disappointment.⁶ It is better to acknowledge that conceptual clarification can only take us so far in understanding a phenomenon than it is to provide a more ambitious account that is unlikely to be extensionally adequate.

2. Political Experiments and Manipulation

Consider the following cases:

- (a) *Informant*. Rose and Sam are siblings who belong to a local church. Sam knows that their grandmother, with whom Rose has a close relationship, is very concerned about Rose's religious well-being. Frequently Rose doesn't manage to get to church, a fact that she does not report to her grandmother, although she also doesn't lie about it. One day, Sam tells Rose that he intends to visit their grandmother in the next month, noting in passing that he suspects Rose's recent church attendance will likely come up in conversation.

⁵ For example, in his *Manipulatory Politics*, Robert Goodin presents a view according to which “analytically, manipulation has two components: it is both deceptive and unwelcome” (Goodin 1980). We do not accept either of these conjuncts. We will soon provide several cases in which some actor attempts to influence another's choice, in a way that at least might be objectionable, without using *deception*—that is, without inducing any false beliefs. Further, we see no principled reason, to suppose that manipulation should be unwelcome. It is psychologically possible that some people will welcome the actions of others even when they could have a moral claim against such actions; even if I am glad that you lied to me to spare my feelings, that lie might still be impermissible.

⁶ The basic problem is that a single non-moralized definition cannot capture the nuance of the moral subject matter it seeks to describe. So such definitions tend to attribute moral errors to some cases that appear innocuous, and to fail to attribute moral errors to intuitively objectionable cases. See Pallikkathayil 2011.

- (b) *Test*. Sarah has declared her intention to help her neighbors move on the weekend. John, a friend of Sarah's, is suspicious that she might be less committed to being a good neighbor than she publicly portrays. Without mentioning any knowledge of her plans, he asks her if she'd like to be play tennis—which he knows she enjoys—at the exactly the time when her neighbors will be moving.
- (c) *Persuasion*. Tom and Andrew are members of their local school board who often find themselves at opposing purposes. Unbeknownst to Andrew, Tom has studied Andrew's positions and values to create stock of reasons that he thinks will help win Andrew over to his view. Tom's own perspective is so different that he regards these reasons and values as nonsense, but he uses them in conversation with Andrew because he regards them as most likely to help Andrew come around to his view.

Each of these cases describe instances where one person intentionally attempts to exercise influence over another, and does so in a way that might at least raise moral concerns. Though they are different circumstances, the common thought is that one agent might be manipulating another. In saying this, we are not yet registering any judgments about the actions' permissibility or impermissibility, justice or injustice, or the like. Rather, we are observing that the cases call out for further moral consideration. They are *morally salient*.

Parallel concerns might be raised about experimental research and campaign methods.

Consider:

- (a') *Disclosure*. In a well-known experiment, researchers sent voters a flyer entitled, "What if your neighbors knew whether you voted?" The letter then gave information about whether the recipient, as well as other people in the recipient's neighborhood, had voted. The letter described voting as a civic duty, and said that the authors were "taking a new approach" to the problem of non-voting by publicizing who does and does not vote (Gerber, Green, and Larimer 2008).
- (b') *Request*. A recent field experiment tested discrimination among state legislators by sending them letters requesting help registering to vote (Butler and Broockman 2011). Some legislators received a letter from an "alias" from "Jake," a name suggesting the sender was white. Other legislators randomly received a letter from an "alias" from "DeShawn," a name suggesting the sender was black. Neither "Jake" nor "DeShawn" was a real person—the names were chosen due to their racial correlates. The researchers found racial discrimination among state legislators.
- (c') *Persuasion*. Modern campaigns use experiments to determine which arguments will be most persuasive to particular voters, and then deploy arguments in the expected optimal way. Campaign staff do not communicate why they have chosen the reasons they are using, instead

presenting the reasons directly to those voters for whom they expect the arguments will be most likely to persuade. Campaigners may produce reasons for action which *they themselves* believe are bad.

None of the actions in (a')-(c') is obviously morally wrong. But like (a)-(c), each one raises moral considerations that recommend further theoretical reflection, and for similar reasons. In each case, the experimenter attempts to intentionally influence an individual's behavior in ways that might be objectionable. As before, the details differ (in ways that we will consider later), but the pre-theoretical worry is again one of manipulation. It appears that the researcher or campaign staffer is trying to get individuals to do something—not with coercion or overt deception—but with tools that change the individual's practical reasons or deliberative circumstances.

Field experiments in political settings like (a')-(c') are morally salient because they are intentional attempts to influence others' choices. Given the importance of individual autonomy, any intentional influencing of another person's choice will carry moral questions in tow. Are they impermissible? Are they unjust? To repeat—these questions need not have positive answers; both offers and sincere persuasive speech are modes of attempting to influence others' choices, and they are typically permissible. But they *might* be answered positively; lying, deception, and coercion, after all, are typically impermissible. Field experiments like (a')-(c') raise interesting moral questions because we cannot readily place them near either moral pole.

Such moral questions seem to have heightened stakes in the democratic context. All else being equal, paradigmatically wrong influencings seem to be more seriously wrong when they are made upon democratic citizens, as those citizens discharge their political responsibilities. Compare:

Grocery Store: Al asks Bert for directions to the nearest grocery store. Bert, who wants to buy the store's last birthday cake for his son, sends Al in the wrong direction.

Voting: Al asks Bert for directions to the nearest polling booth. Bert, who wants Al's preferred candidate to lose, sends Al in the wrong direction.

Both of these misdirections are wrong—but the second, democratic one seems to be more seriously so.⁷ This implies that we have additional reason to investigate the moral considerations that bear upon the permissibility of field experiments (a')-(c'), over and above our reason to investigate the moral considerations that bear upon their interpersonal analogues (a)-(c). Just as deception is more seriously wrong in the democratic context, so too, we suspect, is manipulation. Though it is unclear whether certain field experiments *are* manipulative, we should be especially concerned to find out whether they are.

Having now introduced some actual political experiments, we should clarify how we selected these cases before concluding this section. Throughout this paper, we look at particular experiments because these studies raise the kinds of questions of interest to us. So we emphasize that we are, of course, not entitled to any generalizations about the actual universe of experiments within political science. It is entirely possible that our selections are unrepresentative, and it altogether possible (and perhaps likely) that most experiments will not raise any of the concerns we discuss. Our interest is in the normative issues associated with political experiments, not in empirical generalization about the presence or absence of those issues. From a moral point of view, we could equally well consider these questions exclusively with hypothetical cases. But it is useful to discuss actual experiments, as possible, to illustrate that these issues do—however occasionally—arise.

⁷ This case is just a test is just a test of intuitions. It is not an argument from which the greater relative seriousness of the democratic condition can be *inferred*. Moreover, not all readers will share this intuition. The contrast in these cases is aimed only to prompt reflection on whether the democratic context presents distinctive normative concerns. Whether this inference is vindicated depends on the arguments to follow.

3. The Instrumental Account

So far, we have been canvassing intuitions to support two claims: first, field experiments are worthy of moral investigation, even when they are neither coercive nor deceptive; second, whenever manipulation is morally wrong, it is likely to be *more seriously wrong* in the democratic context. It is high time to see if these intuitions can be explained by more general moral principles.

Randomized experiments can seem to introduce a hostile species into a political habitat. Democratic politics, in one way or another, aims to give its participants an equal say in shaping laws and institutions. Actual practice falls far short. Still, if we take this as a worthy ideal, field experimentation can seem to undermine our democratic ideal. What makes experiments the gold standard for causal inference is precisely what makes them democratically worrisome. They don't just "ask the world a question" (Paul and Healy 2016). They selectively treat a subset of citizens, and leave others untreated, randomizing them into treatment and control conditions. In other words, an experiment, to earn its name, must aim to be causally momentous. If experiments must have the possibility of being consequential, there is the chronic possibility that it will affect the outcome of an election. And this is at least cause for concern, from a democratic point of view.

To see whether this concern is justified, we first need to identify just *what* the concerning fact is. In this section, we will examine one idea, floated by Alan Gerber, according to which field experiments are objectionable only if they actually *swing* elections—that is, when their interventions are but-for causes of electoral outcomes. We believe that this claim is false, and revealingly so. Interventions in elections can be wrong on *non-instrumental* grounds—that is, quite apart from the effects that they produce. In the section that follows

this one, we will develop such a non-instrumental account of why field experiments might be objectionable, on democratic grounds. This section helps to clear the way for our account.

Gerber writes:

[T]here are countless other mundane and essentially arbitrary contributions to the outcome with electoral consequences that are orders of magnitude larger than the typical experimental intervention. A partial list includes: ballot order (Miller and Krosnick 1998), place of voting (Berger, Meredith, and Wheeler 2008), the number of polling places (Brady and McNulty 2004), use of optical scan versus punch card ballots (Ansolabehere and Stewart 2005), droughts, floods, or recent shark attacks (Achen and Bartels 2004), rain on election day (Knack 1994), and a win by the local football team on the weekend prior to the election (Healy, Malhotra, and Mo 2009). That numerous trivial or even ridiculous factors might swing an election seems at first galling, but note that these factors only matter when the electorate is very evenly divided. In this special case, however, regardless of the election outcome, an approximately equal number of citizens will be pleased and disappointed with the result. As long as there is no regular bias in which side gets the benefit of chance, there may be little reason for concern (Gerber 2011, p. 130).

Gerber appears to be making three arguments. First, field experiments are very unlikely to swing elections (because they are efficacious “only when the electorate is very evenly divided”). Second, even if experiments *did* swing elections, they would not be objectionable on that basis because arbitrary factors swing elections all the time, without any great moral import. Third, difference-making experiments are unobjectionable because they do not regularly favor one side rather than rather than the other.

Gerber’s first argument rests on an empirical claim, the full assessment of which would require rigorous testing. For the moment, we will simply raise some grounds for skepticism. In tight races, it is very much a live possibility that experiments can swing an election. Indeed, given the hundreds of get-out-the-vote (GOTV) experiments that have been conducted at every level of American politics, it is quite likely that an experiment *already has* swung an election.

We have more to say about Gerber's second and third arguments, which make normative claims. We think that they prove too much — and dangerously so. Gerber's second argument contains a conditional premise: If the electorate is very evenly divided, then an arbitrary cause affecting the outcome is not democratically important. Is this premise true? To investigate it, we might consider an analogue in first-personal action:

Peer Pressure: You are deliberating about whether to attend a department picnic, or work on a paper. In making this decision, you are aware that you are easily affected by small or arbitrary factors: whether it looks sunny outside, whether your colleague in the next office is planning to go, if you happen to feel hungry twenty minutes before the event, and so on. I know that you are on the fence, and easily swayed by arbitrary considerations, but I have a strong interest that you attend the picnic. I stop by your office and suggest that everyone is likely to go (even though I know this is false). Compelled by peer pressure, you accede to my request.

In this case, it seems you would be justified in resenting me for manipulating you. I influence your action by inducing you to adopt a false belief. It does not seem to excuse my action that I know you are apt to act for arbitrary reasons anyway, when you are on the fence about some action. In general, there is good psychological evidence that humans are easily swayed by arbitrary facts (Doris and Stich 2005).⁸ But the complaint with manipulation is not primarily about acting sub-optimally. The objection is to having one's agency interfered with by another agent. It is about how we are *treated*, regardless of whether that treatment is consequential in the world, e.g., leads us to change our vote or to turnout when we otherwise wouldn't. There is no inconsistency in being happy enough to act arbitrarily, but still indignant over being deceived. The same may be true of the democratic case. If there is some value to democratic decision making, it may be wrong to interfere with such collective choices, even if individuals doing the choosing are not voting on especially good reasons.

⁸ Consider the finding that people who have just found a dime are 22 times more likely to help a person who had dropped some papers on the ground than are those who did not find a dime.

Now let's consider Gerber's third argument:

Premise (1). As long as there is no regular bias in which side gets the benefit of chance, there may be little reason for concern.

Premise (2). An election experiment was causally sufficient to alter the outcome of Election A.

But, Premise (3). For all we know, the next causally sufficient intervention will alter the outcome of Election B in favor of the other party.

Conclusion. The intervention in Election A isn't democratically concerning.

Even if we grant premise (1), the conclusion need not follow. There may be one-off cases in which an election among an evenly divided electorate is swung by some external intervention, such that we regard these cases as democratically concerning even if believe a future intervention is just as likely to result in an opposite outcome. Imagine discovering that Russian interference was sufficient to swing the outcome of the 2016 election. We might well regard this as an affront to democratic values, even if we believed that the next interference might just as likely swing an election in an opposing direction. So the argument is invalid.

We've argued that Gerber's account does not capture all that might be democratically objectionable about field experiments. But there is a broader lesson here. Electoral interventions can be wrong, not simply in virtue of whether they affect the results of the election; they can be wrong because they treat citizens in objectionable ways.

4. An Egalitarian Account of Democratic Manipulation

We will now develop an account of what might be democratically objectionable about the treatment of experimental subjects. If certain field experiments are inconsistent with democratic values, that is so (in part) because they are inconsistent with *political equality*. In

other words, they would be objectionable because they threaten the equal standing of co-citizens.⁹

4.1. *Political Equality*

The ideal of political equality will require elaboration. Let us begin by delivering the intuition that some form of equal standing might be valuable. Perhaps the clearest illustration of this idea occurs when one person—and only that person—is denied the opportunity to vote. Many people regard this as a paradigmatic injustice—a wrong *to* that individual, such that she in particular is aggrieved. But why is that so? We cannot explain this injustice in instrumental terms; it is a familiar point that, given the size of modern democracies, each individual’s vote is nearly worthless in terms of what it can achieve for her. A more compelling explanation is that denying this person an opportunity to vote creates some morally objectionable relation between her and other citizens: she is now inferior; they superior. Democracy is valuable not simply because of the results it achieves—wise policies, peaceful transitions of power—but because of the equality it embodies (Beitz 1989; Waldron 2001; Buchanan 2002; Christiano 2008; Stiliz 2011; Kolodny 2014a; 2014b; Viehoff 2014).

At a second pass, we can distinguish between two sorts of demands that political equality might make upon political arrangements. First, it might make demands upon how power is distributed between ordinary citizens and elite political officials (such as

⁹ This basic idea has several antecedents—both among philosophers, several of whom have claimed that interpersonal relations are of paramount moral importance (Rawls [1993] 1996; Scanlon 2008; Scheffler 2010; Gutman and Thompson 1996); and in the *Belmont Report* (1978) which explicitly brings questions of justice to bear on social science research (p. 9).⁹ For an application of the *Belmont Report* to political science, see Gubler and Selway (2016).

legislators, bureaucrats, and judges), if such officials are to exist.¹⁰ Second, it might make demands upon how power is distributed among ordinary citizens.

In this paper, we will mostly set aside the first set of demands (though we will have something to say about it). We will instead concentrate upon the second set, for it is these demands that govern the relations between political scientists and experimental subjects.

When it comes to relations among ordinary citizens, it is commonly thought that political equality demands something like equal power (hence, one person; one vote) (Brighthouse 1996; Fishkin and Luskin 2005; Goodin and Tanasoca 2014; Viehoff 2014; Goldman 2015). We embrace this, and—for sake of concreteness—we shall embrace Niko Kolodny’s specification of this demand, according to which political equality obtains among ordinary citizens just when they enjoy equal opportunity for free, informed influence (Kolodny 2014b, 310). More precisely, this condition is satisfied when each person has equal *a priori* chances of being decisive over a political choice. For each person, and all of their acts of influence, it must be just as likely that [had X acted differently, the political choice would have turned out differently] (Kolodny 2014b).

4.2. *The Input and Output Objections*

We now consider how field experiments might upset political equality, so understood. There are two ways by which they might.

First, they might upset equality between experimenters and *their subjects*. If an experimenter manipulates a subject into acting, her action is not *freely* made. When experiments intervene in elections, these actions are acts of political influence. Thus, if an

¹⁰ Of course, whether one *is* an elite official is constituted by the political powers that one has. We don’t mean to suggest that the concept of an elite official can be specified independently of political powers.

experimenter manipulates a subject to vote, or to vote in a certain way, the experimenter has a greater opportunity for free, informed influence than the subject does. Call this the *input objection*. The site of its concern is the process of reasoning that precedes the vote. Is it objectionably distorted?

Second, field experiments might upset equality between experimenters and *third-parties*. Not everyone has the knowledge or the resources to mount a field experiment. Thus, when an experimenter influences a political process by those means, she is exercising an opportunity for influence that is not equally shared. Call this the *output objection*. It points not to the relation between experimenters and subjects, but rather between experimenters and everyone else. We're moved by this ideal when we worry about Facebook determining an election, by placing just the right messages in millions of news feeds and making Democrats turn out in droves but leaving Republicans cold. Due to their knowledge and their ownership of a hugely important social network, the owners of Facebook are able to influence politics in a way that few (if any) others can. And this might seem to strike against political equality, just as surely as if the owners of Facebook decided to parley their knowledge and resources by *bribing* the exact right voters.

4.3. *The Crucial Issue: Could Field Experiments Undermine Freedom?*

Both egalitarian concerns about political experiments rest on a claim about freedom: namely, that those experiments subvert the free choice to vote. Freedom is obviously relevant to the input objection. But it is also relevant to the output objection, for reasons we will now explain.

The unequal distribution of means for influencing other citizens need not translate into political inequality (Brighouse 1995; 1996; Kolodny 2014b). Consider, for example, Paul

Krugman. In virtue of his *New York Times* column, Krugman has a vastly greater ability to influence elections than does the typical baker. But this, in itself, does not seem to threaten political equality. The existence of Krugman's column does not seem objectionable in the same way as if Krugman coerced, deceived, or bribed voters (Brighthouse 1995; 1996; Kolodny 2014b).

How can we distinguish between writing a column for the *New York Times* and buying votes? The distinction seems to be founded on the *mode* by which one person attempts to influence another (Brighthouse 1995; 1996; Kolodny 2014b). Some interventions don't distort citizens' thought processes objectionably; others do. And, *for this reason*, unequal distributions of means of influence sometimes do not translate into political inequality, and sometimes do. The output objection, then, also presupposes that field experiments undermine the freedom of experimental subjects.

To determine whether the political egalitarian concern about field experiments has merit, we would need to rigorously examine this presupposition. Unfortunately, the most we can do here is to call for further research. To our minds, drawing a principled line across various methods of influence is a pressing task for political egalitarians—one they have not successfully discharged (Brighthouse 1995; 1996; Kolodny 2014b), and one which we cannot discharge here.

Instead, we note the variety of ways in which freedom might be lost. First, and most obviously, freedom is subverted when eligible options are removed or burdened (Carter 1999). If an experimental treatment imposes a cost, for example, to non-voting that is not otherwise present, then that option is burdened. Second, freedom can be subverted by constraining opportunities for acting with others. Part of the importance of relationships of

equality is that they enable genuinely shared agency (Shiffrin 2008). Agents in unequal relationships cannot reliably count on others' future collaboration, which compromises their ability to engage in long term planning. Where planning opportunities are limited, future choices are constrained. Third, even if an agent's choice is not interfered with, the prospect of possible interference can, itself, limit freedom (Pettit 1997; List and Valentini 2016). Fourth, freedom can be subverted by impairing an agent's control over her own capacities. On some views of personal autonomy, manipulating an agent's sense of self might have this effect by causing the agent to be guided by one self-aspect rather than another (See Mills 1995; Barnhill 2014).¹¹ We are not claiming that *all* of these varieties of freedom are valuable, let alone valuable to the same degree or in the same ways. Our point, in mentioning them, is to suggest the breadth of possible concerns that bear upon the issue of freedom.

5. Implications of the Theory

So far, we have articulated a possible tension between political equality within a democracy and some uses of field experiments. We should be very clear about what we have tried to do. We have not shown that such experiments are, therefore, morally objectionable. In fact, we recognize that the democratic egalitarian's presuppositions are controversial in various ways. Some deny that political equality is non-instrumentally valuable (Dworkin 2000; Arneson 2004; Brennan 2016); others believe that political equality, properly understood, does not require strict equality of opportunity for influence (Beitz 1989; Christiano 1996; Pevnick 2016). Meanwhile, as we have pointed out, it's unclear whether

¹¹ Consider, for example, an experiment that influences citizens by asking them how important it is "to be a voter" rather than "to vote" (Bryan, Walton, Rogers, and Dweck 2011). We grant that this intervention initially seems innocuous. Our point is that it is one sort of intervention that a fuller theory of democratic manipulation should consider.

experimental interventions undermine subjects' free choices to act politically. If they do not, then, for reasons we have outlined, it is difficult to see how they undermine political equality. We have attempted to show, first, that the democratic complaint is worthy of consideration; and, second, to illuminate what one must think about in order to determine whether the complaint can be vindicated.

However these issues are to be resolved, though, one might think that the practical upshot of our discussion is an extreme position. If the objection *is* vindicated—if it true that political equality matters and that field experiments undermine it—then it might seem that *all* experiments in politics are impermissible. Call this rigorist view *prohibitionism*. On the other hand, if the objection is *not* vindicated—if either of its premises is false—then it might seem that all field experiments are permissible (conditioned on meeting the usual requirement of social-scientific experimentation). Thus there would be nothing distinctive about an experiment in a market setting versus an explicitly political setting; experiments are unremarkable forms of free speech, no more objectionable than the bulk buying of ads by a SuperPAC. Call this view *permissivism*.

As far as we know, no one has explicitly defended either of these views. Yet many find themselves pulled towards one or the other.¹² We find each of them uncomfortable in its own way. In the rest of this section, we will articulate a stable position that avoids either extreme. It acknowledges the epistemic—and perhaps democratic—potential of campaign experiments without denying the potential force of the egalitarian objection.

5.1 Which Treatments Are of Egalitarian Concern?

¹² Some practitioners, however, tend towards one or the other. Teele (2014) tends towards prohibitionism; Gerber (2011) and Levitt and List (2009) tend towards permissivism.

One virtue of our egalitarian account is that it has the resources to distinguish among different kinds of experimental interventions. Even if the egalitarian account's premises are true, it follows that only certain interventions are objectionable on egalitarian grounds. These are the interventions that undermine subjects' ability to freely influence political decisions.

Thus, our account delivers the intuitively correct judgments about certain interventions which appear to be permissible. Consider an example from Gerber and Green:

Suppose that a political scientist were to run for public office. Suppose that she made ten or twenty speeches in attempt to curry political support. We ordinarily would not question the rectitude of this behavior even if the churches were chosen at random. Now suppose that the political scientist were interested not simply in winning the election but also in finding out how much these speeches affected the election outcome (Green/Gerber 2002, p. 830).

Gerber and Green believe, and we agree, that there is no difference between the original case and its variation. If an intervention is permissible, it does not become impermissible simply because the intervening agent gains some knowledge from it. Our account vindicates that judgment about this particular case, because both interventions—involving political speeches—influence voters through a paradigmatically freedom-preserving means.¹³

The moral character of other interventions is less clear. Section 2 introduced a recent experiment that came to prominence in part for its controversial treatment of disclosing voting records (Gerber, Green, and Larimer 2008). Researchers gathered turnout information for 180,000 voters, dividing them into a control group and four treatment groups. Each treatment group received different type of encouragement to vote in an

¹³ More generally, our account does not imply that interventions become objectionable simply because information is gained from them, for it does not hold that the *information gain* is the putative wrong-making feature of experiments.

upcoming election: (1) a simple reminder to vote, claiming that voting is a civic duty; (2) the same reminder, except with the added information that researchers were watching the election to learn about turnout; (3) the reminder as in (1), with an added note that turnout is a matter of public record; and (4) a letter displaying the turnout for the recipient and the recipient's neighbors in a previous election. Treatments (1)-(3) raised participation from between 1.8 and 4.9 percent relative to the control group. Treatment (4) raised participation dramatically—8.1 percent relative to the control group.

The researchers developed their experimental design after Mark Grebner, a political consultant, self-funded an experiment showing that disclosing voting history was dramatically more effective than other GOTV mailers. Grebner had observed large numbers of voters misreported having voted when in fact they had not. He came to “suspect that a substantial number of potential voters believed that voting is essentially a private act, rather than a public one, and it was safe to routinely misreport voting not only to researchers but also to friends and family” (qtd. in Green and Gerber 2010, p. 332). Thus, his strategy was to use the technique of publicizing voting “by making them think they couldn’t continue to get away with it.” As the authors describe, there are strong dispositions toward compliance with social norms, so the prospect of “others believed to be watching” motivates actions to “avoid shame and social ostracism” (Gerber, Green, and Larimer 2008).

This study quickly proved controversial, with some recipients contacting local District Attorneys and the researcher being contacted by law enforcement (Michelson and Nickerson 2011, 439). Our conception of manipulation can help to appreciate the intensity of their concerns. The researchers disclosed information that would otherwise would have been very unlikely to be accessed, and they did so with the explicit intent of creating a penalty

of moralized shame that would influence the actions of other persons. *Prima facie*, the researchers' action looks like a case of manipulation. The shaming threat impinged upon subjects' ability to freely influence political processes.

The authors might counter—as in fact their flier to voters declared in large type font—that voting records are public information. But it is unclear why this should be probative. In the case of Informant (from section 1), Sam threatens to visit Rose's grandmother and reveal her lack of church attendance. The case might be stylized to make clear that her non-attendance is public in the relevant sense—anyone could find out by asking her fellow parishioners (among whom it is not a secret), attending the publicly open church one Sunday themselves, and so on. What Sam does is threaten to deploy this information in a way that is damaging to Rose, by compromising her relationship with her grandmother. Likewise, the researchers threaten to take information which is already public but non-damaging, and use it in a way which will potentially compromise some citizens' relationships with other citizens. Indeed, this is precisely the researchers' intention, since it is the threat of shame among fellow citizens (neighbors) through which theorized treatment mechanism works.

Still other studies seem (again, provided that the basis of the egalitarian concern is solid) more likely to be objectionable on egalitarian grounds. Consider one combining elements of both shame and deception. A recent experiment investigated different strategies for increasing voter turnout in three municipal elections (Ely, Iowa; Monticello, Iowa; and Holland, Michigan) (Panagopoulos 2010). In two of these communities, voters received a mailer that those who voted in the election would be honored by having their names appear in a local newspaper. In the third community, citizens received a mailer that those who did

not vote in the upcoming election would be shamed by having their names revealed in a local newspaper. The pride treatment increased turnout by .8-1.5% and by 4.5-4.7%. The shame treatment increased turnout by 6.3-6.9 percent. The researcher found that in general, shame is more effective than pride, because shame motivates both high- and low-propensity voters, while pride motivates high-propensity voters only (Panagopoulos 2010, p. 382)

We are uncertain of whether the threat to shame was not actually carried out. In one sense, not carrying out the threat might be thought to be morally preferable. However, failing to publish the names also makes the action more clearly manipulative. If this is not what happened, let us imagine a related case in which the researcher does not carry out the threat. Under these stipulated conditions, the researcher influenced the actions of persons by affecting their beliefs and emotions through making explicitly false statements. This is the sort of case which we would take to be wrong in everyday circumstances, even if the manipulated action were judged to be morally desirable. Suppose that some aid organization keeps records of donors which can be accessed at its website. However, as a member of the organization who is well aware of its acute funding needs, I decide that I can increase donations by publishing the names of every person whom I ask for a donation, but declines. If I threaten potential donors with exposing their uncharitable refusals, and I then do not actually publish any names, they could reasonably object that I undermined control of their own agency through lying to them.

5.2 A Distinction Between Citizens and Elite Political Officials

Our account also suggests that there might be a distinction between experiments conducted on ordinary citizens and those conducted on elite political officials (such as legislators, bureaucrats, and judges). That is so, because it is possible that political equality

makes different demands upon the treatment of ordinary citizens and the treatment of elite officials. Perhaps certain forms of treatment which would *create* political inequality between ordinary citizens *collapse* the inequalities of power that inhere whenever some citizens wield special political powers in official roles.

To further substantiate this thought, it might help to have an example in mind. Earlier we considered a study which asks, “Do Politicians Racially Discriminate Against Constituents?” (Butler and Broockman 2011). In this experiment, state legislators were sent emails asking for help registering to vote. These emails were sent either under a name typically perceived as black (DeShawn Jackson), or a name typically perceived as white (Jake Mueller). In total, the researchers sent 4,859 emails to state legislators, and received 2,747 responses, finding that “Democratic and Republican whites [legislators] discriminated against the black alias at nearly identical and significant levels (by 6.8 percentage points and 7.6 percentage points, respectively)” (Butler and Broockman 2011, p. 472). In this case, the researchers influenced the actions of state legislators by causing them to form a false belief—namely, that there was some particular person who had both sent the request and needed help registering to vote. Although the researchers’ emails did not contain any actual false assertions, and so do not qualify as lying, this fact does not seem to eliminate the suspicion that their intervention involved some objectionable form of influence. After all, it would hardly have served the deceptive purpose of the request to include a sentence of the form, “And I, Jake Mueller, am an actual person and constituent!” (cf. Langton 1992).

Our account offers some reasons for thinking that this experiment (performed upon legislators) might be *disanalogous* to a variation performed upon ordinary citizens. That is so, because the experiment may not have created or aggravated any objectionable

inequalities of influence. First, as we noted earlier, political equality need not require strictly equal opportunities for influence between both ordinary citizens and elite officials. Thus, even if this experiment *did* create inequalities of influence, those inequalities need not have been *objectionable*. Second, the experiment likely did *not* create or aggravate inequalities of influence; if anything, it lessened them. The experiment uncovered important information that could be used to hold relatively powerful legislators accountable to ordinary citizens. Thus, it helped to reduce an extant inequality. Furthermore, given the threats to social equality posed by racist attitudes—particularly among elites—the researchers create a way of potentially exposing these attitudes for critical scrutiny. This may help advance equality among citizens, while not creating any problematic inequality between the researchers and ordinary citizens.

5.3 Do Epistemic Benefits Matter?

Finally, we address the moral importance of the information yielded by field experiments. Some social scientists have been tempted by the idea that the epistemic gains of research might justify treatment of subjects (see Levitt and List 2009). There are three ways in which it might.

First, as the Butler and Brookman study shows, the acquired knowledge might help to *eliminate* extent political inequalities. In that case, the epistemic gain from a field experiment does not merely *override* the egalitarian objection to that experiment; rather, it *undercuts* that objection, leaving it without any normative weight.¹⁴

¹⁴ An undercutting defeater is some consideration in the presence of which what would have been a reason is no longer a reason (Schroeder 2011, p. 334).

Second, the acquired knowledge might have intrinsic value. Knowledge of political behavior, just like Einstein's field equations and the axiom of choice, is just a good thing to have—entirely apart from its usefulness. If this kind of value makes certain field experiments permissible, in spite of the egalitarian objection, it would have to do so by *overriding* the egalitarian objection. In other words, the knowledge acquired by the experiment would have to be of greater moral significance than the political inequalities caused by it.

We are skeptical of this possibility. We set aside whether the sort of knowledge acquired by field experiments has the same intrinsic significance as, say, Einstein's field equations. The more pressing problem is that political equality seems to be a rather *stringent* value—that is, in order for actions inconsistent with its demands to be permissible, they must realize some value of great moral significance.¹⁵ The intrinsic value of the knowledge yielded by field experiments seems not to rise to this level. Consider, by way of analogy, a campaign that forcibly suppresses voters who would have voted down funding for political science departments at public universities. Though this campaign would, as a foreseeable consequence, increase our stock of social scientific truths, it hardly seems morally permissible. Of course, the analogy is imperfect—out and out suppression may undermine political equality to a greater degree than the interventions effected by field experiments. Accordingly, the moral calculus may be different in the two different phenomena. But the analogy does provide some grounds for doubting that field experiments would be permissible, even if the egalitarian objection is sound, simply because they yield intrinsically valuable knowledge.

¹⁵ For this notion of stringency, see Tadros 2016, 102.

Third, the acquired knowledge might have non-egalitarian instrumental value—that is, it might be valuable because useful, but not as a means of reducing political inequalities. Perhaps, for example, the knowledge acquired by a field experiment helps us to enact various policies which increase voter turnout, which in turn improves the substantive quality of our laws.

Structurally, this consideration ought to be treated in much the same way as the intrinsic value of the knowledge gained. Similarly, this form of instrumental value would have to *override* (and not undercut) the force of the egalitarian objection. And, similarly, it would have to do so in spite of the fact that political equality seems to be a rather stringent value. However, one difference might be that improvements to substantive justice frequently will be of greater moral significance than the intrinsic value of knowledge. Thus, it is more likely that experiments which genuinely do yield such practically valuable knowledge would be permissible, even if the egalitarian objection is sound.

Conclusion

In their handbook for non-academics, the pioneers of political experiments insist that “experimentation is not some special form of sorcery known only to Yale professors. Anyone can do it...” (Green and Gerber 2008, p. 11). The objections that we’ve raised here fall far short of the charge of sorcery. But if what makes sorcery objectionable, in part, is its ability to manipulate agents, then a subset of experimental manipulations in campaigns may invite the same prohibition as sorcery. At the same time, it’s too easy to support the two absolutist views that we began with: all-out prohibition or all-out permission. The simplicity of each of these views, however appealing, isn’t enough to motivate them. What they neglect is the

special setting of elections – an environment whose value is underwritten by an ideal of democracy. We've suggested that experiments in political settings must reckon with distinctively democratic worries – the value we place in having an equal say, and the way experiments can threaten that complex web of egalitarian relationships. A theory of democratic manipulation, developed further, would help us see what is objectionable about certain experimental interventions, and what makes others perfectly compatible with our having an equal say.

Bibliography

- Achen, Christopher H., and Larry M. Bartels. "Blind retrospection: Electoral responses to Drought, Flu, and Shark Attacks." (2004).
- Ansolabehere, Stephen, and Charles Stewart Stewart. "Residual votes attributable to technology." *Journal of Politics* 67.2 (2005): 365-389.
- Arneson, R. 2004. "Democracy is not Intrinsically Just." In *Justice and Democracy: Essays for Brian Barry*, Keith Dowding, Robert Goodin, and Carole Patemen eds., 40 - 58. Cambridge: Cambridge University Press.
- Barrett, C.B. and M.R. Carter. 2014. "A Retreat from Radical Skepticism: Rebalancing Theory, Observational Data, and Randomization in Development Economics," in *Field Experiments and their Critics*, 78-114.
- Baron, M. 2003. "Manipulativeness." *Proceedings and Addresses of the American Philosophical Association* 77(2): 37-54.
- Barnhill, A. 2014. "What is Manipulation?" in C. Coons and M. Weber [eds.] *Manipulation: Theory and Practice*. Oxford: Oxford University Press.
- Beitz, C. 1989. *Political Equality*. Princeton: Princeton University Press.
- Brennan, J. 2016. *Against Democracy*. Princeton: Princeton University Press.
- Bryan, C.J., Walton, G.M., Rogers, T., and C.S. Dweck. 2011. "Motivating Voter Turnout by Invoking the Self." *PNAS* 108(31): 12653-12656.
- Brighouse, H. 1995. "Political Equality and the Funding of Political Speech." *Social Theory and Practice* 21 (3): 473-500.
- . 1996. "Egalitarianism and Equal Availability of Political Influence." *Journal of Political Philosophy* 4 (2): 118-141.
- Buss, S. 2005. "Valuing Autonomy and Respecting Persons: Manipulation, Seduction, and the Basis of Moral Constraints," *Ethics* 115: 195-235.
- Butler, D. M. and D. E. Broockman. 2011. "Do Politicians Racially Discriminate Against Constituents? A Field Experiment on State Legislators." *American Journal of Political Science* 55(3): 463-477.
- Carter, I. 1999. *A Measure of Freedom*. Oxford: Oxford University Press.
- Cave, E. 2007. "What's Wrong with Manipulation?" *Ethical Theory and Moral Practice* 10: 129-144.
- Conly, S. 2004. "Seduction, Rape, and Coercion." *Ethics* 115: 96-121.
- Coons, C. and M. Weber. 2014. *Manipulation: Theory and Practice*. Oxford: Oxford University Press.
- Desposato, S. 2016. *Ethics and Experiments: Problems and Solutions for Social Scientists and Policy Professionals*. New York: Routledge.
- Doris, J. M. and S. P. Stich. 2005. "As a Matter of Fact: Empirical Perspectives on Ethics," in F. Jackson and M. Smith [eds.] *The Oxford Handbook of Contemporary Analytic Philosophy*. Oxford: Oxford University Press.
- Druckman, J.N., D.P. Green, J.H. Kuklinski, and A. Lupia. 2011. *Cambridge Handbook of Experimental Political Science*. Cambridge: Cambridge University Press.
- Dworkin, R. 2000. *Sovereign Virtue*. Cambridge: Harvard University Press.
- Fishkin, James S., and Robert C. Luskin. 2005. "Experimenting with a Democratic Ideal: Deliberative Polling and Public Opinion." *Acta Politica* 40:284-98.
- Gerber, A. S., D. P. Green, and C. W. Larimer. 2008. "Social Pressure and Voter Turnout: Evidence from a Large-Scale Field Experiment." *American Political Science Review* 102(1): 33-48.

- Gerber, A.S. 2011. "Field Experiments in Political Science," in Druckman et. al.
- Green, D.P. and A.S. Gerber. 2008. *Get Out the Vote: How to Increase Voter Turnout*. 2nd ed. Washington, D.C.: Brookings Institution Press.
- Donald P. Green and Alan S. Gerber, "Reclaiming the Experimental Tradition in Political Science," *Political Science: The State of the Discipline*, ed. Ira Katznelson and Helen V. Milner (New York: W.W. Norton & Company, 2002), pp. 805-832,
- Goldman, Alvin I. 2015. "What is Democracy? (and What is its Raison D'Etire)?" *Journal of the American Philosophical Association* 1 (2): 233-256.
- Greenspan, P. 2003. "The Problem with Manipulation." *American Philosophical Quarterly* 40(2): 155-164.
- Godfrey-Smith, P. 2012. "Metaphysics and the Philosophical Imagination." *Philosophical Studies* 160(1): 97-113.
- Goodin, Robert. 1980. *Manipulatory Politics*. New Haven: Yale University Press.
- Goodin, R. and A. Tanasoca. 2014. "Double Voting." *Australasian Journal of Philosophy* 92(4): 743-758.
- Green, D. P. and A. S. Gerber. 2010. "Introduction to Social Pressure and Voting: New Experimental Evidence," *Political Behavior* 32: 331-336.
- Gubler, J.R. and J.S. Selway. 2016. "Considering the Political Consequences of Comparative Politics Experiments," in *Ethics and Experiments*, pp. 171-182.
- Gutman, A. and D. Thompson. 1996. *Democracy and Disagreement* (Cambridge: Harvard University Press, 1996).
- Herman, B. 1993. *The Practice of Moral Judgment*. (Cambridge: Harvard University Press, 1993).
- Hill, T. 1991. "The Importance of Autonomy," in *Autonomy and Self-Respect*. Cambridge: Cambridge University Press.
- Healy, K and Paul, L.A. 2016. Transformative Treatments. *Nous* 10.111/nous.12180.
- Kolodny, N. 2014a. "Rule Over None I: What Justifies Democracy?" *Philosophy & Public Affairs* 42(3): 195-229.
- . 2014b. "Rule Over None II: What Justifies Democracy?" *Philosophy & Public Affairs* 42(4): 287-336.
- Langton, R. 1992. "Duty and Desolation." *Philosophy* 67: 481-505.
- Levitt, S.D., and J. A. List. 2009. "Field Experiments in Economics: The Past, the Present, and the Future," *European Economic Review* 53:1, 1-18.
- List, C. and L. Valentini. 2016. "Freedom as Independence," *Ethics* 126: 1043-1074.
- Manne, K. 2014. "Non-Machiavellian Manipulation and the Opacity of Motive," in C. Coons and M. Weber [eds.] *Manipulation: Theory and Practice*. Oxford: Oxford University Press.
- Michelson, M. R. and D. W. Nickerson. 2011. "Voter Mobilization," in J. N. Druckman, D. P. Green, J. H. Kuklinski, and A. Lupia [eds.] *Cambridge Handbook of Experimental Political Science*. Cambridge: Cambridge University Press.
- Miller, Joanne M. and Jon A. Krosnick. 1998. "The Impact of Candidate Name Order on Election Outcomes." *Public Opinion Quarterly* 62: 291-330.
- Mills, C. 1995. "Politics and Manipulation." *Social Theory and Practice* 21(1): 97-112.
- Pallikkathayil, Japa. 2011. "The Possibility of Choice: Three Accounts of the Problem of Coercion," *Philosopher's Imprint* 11(16): 1-20.
- Panagopoulos, Costas. 2010. "Affect, Social Pressure and Prosocial Motivation: Field Experimental Evidence of the Mobilizing Effects of Pride, Shame, and Publicizing Voting Behavior." *Political Behavior* 32 (2010: 369-386.

- Pettit, P. 1997. *Republicanism: A Theory of Freedom and Government*. Oxford: Oxford University Press.
- Pevnick, R. 2016a. "Does the Egalitarian Rationale for Campaign Finance Reform Succeed?" *Philosophy and Public Affairs* 44(1): 46 – 76.
- . 2016b. "The Anatomy of Debate about Campaign Finance." *Journal of Politics* 78(4): 1184-1195.
- Rawls, J. (1951) 1999b. "Outline of a Decision Procedure for Ethics." *Philosophical Review* 60: 177–97. Reprinted in Rawls 1999b, 1–19.
- . (1971) 1999a *A Theory of Justice*, revised edition. Cambridge: Harvard University Press.
- . (1974) 1999b. "The Independence of Moral Theory." *Proceedings and Addresses of the American Philosophical Association* 48: 5-22. Reprinted in Rawls 1999b, 286-302.
- . (1980) 1999b. "Kantian Constructivism in Moral Theory." *Journal of Philosophy* 77: 515-72. Reprinted in Rawls 1999b, 303-358.
- . (1993) 1996. *Political Liberalism*. New York: Columbia University Press.
- . 1999b. *Collected Papers*, ed. Samuel Freeman. Cambridge: Harvard University Press.
- Rose, J. 2016. *Free Time*. Princeton: Princeton University Press.
- Scanlon, T. M. *What We Owe to Each Other*. Cambridge: Harvard University Press.
- . 2008. *Moral Dimensions*. Cambridge: Harvard University Press. 1998.
- Scheffler, S. 2010. *Equality and Tradition: Questions of Value in Moral and Political Theory*. Oxford: Oxford University Press.
- Schroeder, M. 2011. "Holism, Weight, and Undercutting," *Nous* 45(2): 328-344.
- Shiffrin, S. V. 2000. "Paternalism, Unconcionability Doctrine, and Accommodation." *Philosophy & Public Affairs* 29(3): 205-250.
- Stilz, A. 2011. "Collective Responsibility and the State." *Journal of Political Philosophy* 19(2): 190-208.
- Tadros, V. 2015. "Permissibility in a World of Wrongdoing." *Philosophy and Public Affairs* 44(2): 101-132.
- Teele, D. L. 2014. *Field Experiments and their Critics*. New Haven: Yale University Press.
- . 2014. "Reflections on the Ethics of Field Experiments," in *Field Experiments and their Critics*.
- Viehoff, D. 2014. "Democratic Equality and Political Authority," *Philosophy & Public Affairs* 42(4): 337-375.
- United States. 1978. *The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research*. Bethesda, Md: The Commission.
- Zimmerman, B. 2016. "Information and Power: Ethical Considerations of Political Information Experiments," in *Ethics and Experiments*, pp. 183-197.