

THE PRACTICALITY OF POLITICAL PHILOSOPHY*

BY JUSTIN WEINBERG

I. INTRODUCTION

Political philosophy has two masters: knowledge and action. It purports to tell us what is true, or at least what to believe, about justice, authority, and other fundamental political ideas. And it purports to tell us what to do, insofar as we are to act in certain ways that promote or are at least consistent with these fundamental political ideas. Accordingly, we take a political philosophy to be defective when we think it is mistaken in its account of fundamental political ideas, and also when we think we cannot bring about what it says is just and good or act in accordance with it. In other words, products of political philosophy—the kind I will focus on most in this paper are principles of justice—should be true *and* practical, and are defective when they are either false or impractical.

That, at least, is the commonsense view of political philosophy, if there is such a thing. It is a view captured in John Rawls's term, "realistic utopia," which evokes a concern with both practicality and ideals.¹ Over the past several years, the commonsense view has come under pressure from two directions. From one direction are those who, it could be said, are more interested in "utopia," that is, in figuring out what the best, most just society looks like. Their aim is to create or discover a theory that matches up with an ideal of justice, and they are critical of the notion that theories of justice should deviate from that ideal for the sake of being in some sense realistic. Call these people the "idealists."²

From the other direction are those who are more interested in being "realistic." For them, feasibility is crucial, and they are critical of theories of justice that require particular basic principles of justice or institutional arrangements regardless of whether such principles or arrangements could

* I would like to thank David Schmitz, Carmen Pavel, Rosie Johnson, and several anonymous referees for helpful written comments on an earlier draft of this essay. I am also grateful to the other contributors to this volume for their critical comments and discussion of these ideas. A version of the essay was delivered at the Midsouth Philosophy Conference in Memphis, Tennessee in February of 2013, and I would like to thank my audience there and my commentator, Ashley Acosta-Fox, for their questions and comments.

¹ John Rawls, *The Law of Peoples* (Cambridge, MA: Harvard University Press, 2001), 7. Also see the discussion of "optimization" as the "confrontation of desirability with feasibility considerations" in Geoffrey Brennan, "Feasibility in Optimizing Ethics," *Social Philosophy and Policy* 30, nos. 1/2 (2013): 314.

² In Section III, I explain how G. A. Cohen and David Estlund are idealists in this sense.

be implemented, or whether, once implemented, they would have the intended effects. Call these people the “implementers.”³

While idealists and implementers can be seen as endpoints on a shared spectrum,⁴ and while some individual political philosophers may see themselves at points somewhere between these endpoints, drawing the contrast between the two approaches more starkly can help us better evaluate them and understand their disagreement. We can do this by framing the disagreement between them as over this question: *must principles of justice be practical?* The implementers say yes. The idealists say no.

There are, of course, different concepts of practicality, and it is worth specifying which this question asks about. We can distinguish between what we can call *formal* practicality and *substantive* practicality. Formal practicality involves a notion of “practical” that connotes action, and that philosophers sometimes contrast with “theoretical” or “descriptive.” A principle of justice is formally practical if it tells some agent(s) to do something. This sense of practicality is not at issue in this paper. Rather, my focus is on substantive practicality.⁵ Substantive practicality involves a notion of “practical” that connotes possibility, or possibility within certain evaluative constraints, or success, and that is sometimes contrasted with “impractical,” “unrealistic,” “unfeasible,” and the like. A principle of justice is substantively practical when it is both formally practical and the actions it prescribes meet some to-be-specified set of conditions about the action’s success. Understood this way, the question as to whether principles of justice must be practical asks whether a necessary condition of an acceptable principle of justice is that it is *substantively* practical.⁶

Related to this question is a broader dispute that pits “ideal theory” against “nonideal theory.” As others have discussed, there are multiple conceptions of these terms, with the result that it is not always clear what the dispute is over, or whether one must choose one side or another.⁷ For the purposes of this paper, I take philosophizing about justice to be “ideal” to the extent that it replaces empirically substantiated information about persons, society, and institutions with certain idealizing assumptions, and then proceeds by determining which principles of justice would be appropriate for a society in which these assumptions held. These assumptions

³ The views of specific implementers are described in Section II.

⁴ For a similar spectrum, see Colin Farrelly, “Justice in Ideal Theory: A Refutation,” *Political Theory* 55, no. 4 (2007): 844–64.

⁵ The reader should assume that subsequent references to practicality are to *substantive* practicality, unless otherwise noted.

⁶ The distinction between formal and substantive practicality is important because it clears up a vagueness that can be confusing. See, for example, the discussion in Section III.

⁷ For some general discussions of the distinction, see Laura Valentini, “Ideal vs. Non-Ideal Theory: A Conceptual Map,” *Philosophy Compass* 7, no. 9 (2012): 654–64; A. John Simmons, “Ideal and Nonideal Theory,” *Philosophy and Public Affairs* 38, no. 1 (2010): 5–36; David Schmidtz, “Nonideal Theory: What It Is and What It Needs To Be,” *Ethics* 121, no. 4 (2011): 772–96.

have included, typically, that: (i) individuals comply with the requirements of justice incumbent upon them in their everyday lives, (ii) individuals actively support with their attitudes and actions the principles of justice and the institutions the principles require, (iii) institutions which satisfy the principles of justice are readily available, and (iv) these institutions function well, not by luck but by design, and are not marred by other serious deficiencies.⁸ Theorizing about justice is nonideal to the extent that these idealizing assumptions are tempered by empirically supported, realistic accounts of individual compliance, attitudes and actions, as well as institutional possibilities and functioning. Thus, the ideal/nonideal categorization is one of degree; theories of justice and their constitutive parts are more or less ideal.

Though it is not a precise mapping, it is a fair generalization to say that the implementers, who hold that principles of justice must be practical, tend to favor theory at the nonideal end of the spectrum. According to them, ideal theory leads to results that are not practical, and thus should be rejected in favor of the more realistic nonideal theory. The idealists, who hold that principles of justice need not be practical, in contrast, tend to favor ideal theory. According to them, facts about a principle's practicality are distinct from facts about its justifiability or truth. Though they may admit that ideal theory may not be practical, they do not take that as a reason to reject it.

In short, the implementers say that ideal theory should be replaced (with a kind of nonideal theory) because it does not meet the standard of practicality; the idealists say that the standard of practicality should be rejected because it mistakenly condemns ideal theory.

Despite this disagreement, both the implementers and the idealists tend to share two points of agreement. First, they both accept the characterization of ideal theory as not essentially (substantively) practical.

Second, while implementers and idealists disagree over whether we should hold principles of justice to a standard of practicality, both sides generally agree on what practicality means, accepting what I will call

IP, or the *implementation-prediction* conception of *practicality*:

A principle or set of principles is "practical" if (a) its implementation by way of a series of events that begins at the relevant status quo and ends at something relevantly similar to what the theory prescribes is both plausible and desirable, and (b) the predictions (both descriptive and evaluative) about its implementation which contribute to its justification are sufficiently accurate.

⁸ My conception of ideal theory is based largely on what I take to be Rawls's view. See John Rawls, *A Theory of Justice*, rev. ed. (Cambridge, MA: Harvard University Press, 1999), 7–8. See also p. 4 on well-ordered society; p. 125, on strict compliance; and pp. 215ff on nonideal theory.

There are two central claims of this paper. The first is that IP is not a conception of practicality that philosophers should strive to meet in crafting principles of justice. The implementers, then, make a mistake in holding that principles of justice should meet the IP standard of practicality. The second claim is that there is an alternative conception of practicality, overlooked in the dispute between idealists and implementers, which is appropriate for principles of justice. I call this the *experimentation-learning* conception of practicality, or EL. The idealists, then, make a mistake, too, in their general rejection of practicality.

Both mistakes, I argue, are rooted in kinds of epistemological overconfidence. The implementers are overconfident about our knowledge of what is possible and what we will come to think about it. The idealists are overconfident about our knowledge of what justice is. I argue (in Section IV) that we have good reason to think that our knowledge of justice, of what options are available to us, and what our thoughts will be under various options, is quite limited, and that further experience is necessary to improve it. The positive contribution of this essay (introduced in Section V) is to point toward a kind of political philosophy that recognizes these limits to our knowledge while providing us with some means by which to overcome them.

I do not aim for this essay to settle the debate between ideal and nonideal theory. However, that debate does put pressure on ideal theorists to answer this question: "If ideal theory is not to be implemented, what is it for?" As we will see, the EL conception of practicality can provide an answer to that question that preserves the idealists' treasured independence from the status quo without brushing aside the messy real world to which implementers demand we pay attention.

In what follows, I will provide some background on the implementers (Section II) and idealists (Section III), describing their stances on the practicality of principles of justice, and explaining their common standard of practicality, IP. Following that, I will argue that neither the idealists' rejection of practicality in general, nor the implementers' embrace of IP, is defensible (Section IV). I will then describe the alternative conception of practicality, EL, arguing that it avoids the defects of IP and gives revised, defensible roles to both nonideal and ideal theory (Section V). I then turn to the notion of experimentation at the heart of the EL model and some objections (Section VI) before concluding (Section VII).

II. THE IMPLEMENTERS

Some people who read Rawls's *A Theory of Justice* complain about how detached from reality it seems. Considering its "Original Position," they wonder: Why should we take seriously the results of a hypothetical discussion amongst ignorant imaginary nonpersons under unrealistic conditions about idealistic principles for a society characterized by general

features that will never obtain in the real world? Though I think it has a good answer, this is a fair question. It is an expression of a concern that can be raised about many different theories of justice: What good is this theory for *us, here and now*? *Us*—the kinds of beings we are, characterized by our “human nature.” *Here*—the kind of world we live in and in which our politics takes place, with its incapacity to bend precisely to our wills. *Now*—the particular problems we currently face, and the distance we are now from the just society depicted in the theory. What can this theory tell us about what to do so that our society can be more just?

This concern has evolved in different ways. One development is the advancement of what can be called the “bad effects criticism.”⁹ This is a criticism leveled against some versions of ideal theory, that they will lead to bad effects if we attempt to implement them. Rawls, who introduced the term, “ideal theory,” has been the subject of much of this criticism.

Some of these bad effects are bad because they fail to achieve the aims of the theory. Colin Farrelly, for example, criticizes Rawls for not having a realistic view of the cost of providing for and defending people’s liberties. Rawls aims for his theory to protect liberties and promote welfare, but prioritizes the former over the latter. Farrelly argues that this prioritization may leave people very badly off in terms of welfare in circumstances in which liberties are costly.¹⁰ Rawls also aims to neutralize the effect that factors that are “arbitrary from a moral point of view,” such as race, have on our life prospects. Charles Mills, however, argues that Rawls’s stipulation that he is identifying principles of justice for a society characterized by “strict compliance” with its fundamental principles and laws would, if applied to real world circumstances, give wholly inadequate treatment to issues of racial justice, which have past failures to comply with justice at their core.¹¹

Even when a theory’s aims are ostensibly achieved, there may be bad effects the theorist did not anticipate. Along these lines, Liam Murphy argues against Rawls’s view that principles of justice apply only to the institutions of the basic structure, and not to individuals in their everyday lives. Murphy’s criticism raises the possibility that all of our institutions successfully comply with Rawlsian principles, yet the prospects of the least well-off are poor because the theory does not call upon individuals in their non-official capacities to assist them.¹² And it may be that the best way to interpret Michael Sandel’s early criticism of Rawls is as a complaint about a distorting kind of idealization that would lead to bad effects. The idealization, for Sandel, would be that the parties to the Original Position could

⁹ Justin Weinberg, “A Little Reality is a Dangerous Thing” (unpublished manuscript).

¹⁰ See Farrelly, “Justice in Ideal Theory: A Refutation.”

¹¹ Charles W. Mills, “Rawls on Race/Race in Rawls,” *Southern Journal of Philosophy* 47, no. 1 (2009): 161–84. Also see Mills, “‘Ideal Theory’ as Ideology,” *Hypatia* 20 (2005): 165–84.

¹² Liam Murphy, “Institutions and the Demands of Justice,” *Philosophy and Public Affairs* 27, no. 4 (1998): 251–91.

discern their good in the absence of knowledge about their communal roles and relationships; the bad effects would be the deterioration or absence of communally sustained goods in a liberal society.¹³

In short, if the normative claims of a theory of justice depend for their plausibility and acceptability on certain assumptions that do not actually obtain, then if we act on those claims, things will be bad and justice will not be served well, or, in some domains, not served at all. Note that the complaint here is not simply that society cannot take an ideal theory of justice and immediately and unproblematically implement it. Rather, it is that the idealizations that go into the construction of the theory of justice make it the case that even the gradual implementation of the theory may be impossible or lead to undesirable outcomes.

I leave aside whether these complaints are well-founded. For the purposes of this paper, the important point is what these kinds of complaints have in common: that they are advanced atop a standard of practicality the implementers hold theories of justice to, and to a specific conception of practicality that fleshes out that standard, namely, the IP standard of practicality.

So, for a theory of justice to be IP-practical, it must be ultimately implementable, and the relevant effects of its implementation must be sufficiently similar to what the theory predicts. Since, according to their critics, ideal theories may not be successfully implementable, or if implementable, may bring about bad effects, some question the usefulness of ideal theory.

On some views, ideal theory has the more modest task of describing a "Paradise Island" without providing any guidance about how to get there.¹⁴ But on others, it is useless. Amartya Sen argues that we should abandon what he calls the "transcendental" approach of ideal theory in favor of comparing practically available alternatives.¹⁵ David Wiens makes a substantially similar point. He writes, "One might argue that we need ideal theory to identify the morally acceptable solutions within the feasible set. But there's no reason to employ the overwrought framework of ideal theory to help us here. Attempting to identify ideally just arrangements and the principles that regulate them is liable to distract from the task of identifying solutions."¹⁶ According to these critics, ideal theory can be an obstacle to identifying and implementing justice.

If we were to give the implementers a slogan, it would be this: Justice cannot be impractical.

¹³ Michael Sandel, *Liberalism and the Limits of Justice* (Cambridge: Cambridge University Press, 1982).

¹⁴ Ingrid Robeyns, "Ideal Theory in Theory and Practice," *Social Theory and Practice* 34 no. 3 (2008): 341–62, at 361.

¹⁵ Amartya Sen, *The Idea of Justice* (Cambridge, MA: Harvard University Press, 2009).

¹⁶ David Wiens, "Prescribing Institutions without Ideal Theory," *Journal of Political Philosophy* 20, no. 1 (2012): 45–70, at 67.

III. THE IDEALISTS

While the implementers reject ideal theory for its failure to meet the standard of practicality, idealists have defended ideal theory from the requirement that it be practical. That a theory of justice is impractical is not a reason to think it is false. As G. A. Cohen puts the point, “sometimes justice is unattainable.”¹⁷ Further, that it is unattainable does not render it any less just. In fact, rendering it more attainable may make it less just. He writes, “philosophers in search of justice should not be content with an expedient compromise.” It is a mistake “to call expedience *justice*.”¹⁸ Cohen uses the term “expedient” here, rather than practical or implementable, but the point remains the same.

Similarly, while David Estlund says that societies “should not implement institutions that people will not be able to bring themselves to comply with,” he adds: “The question is whether that is a constraint on the content of justice.”¹⁹ His answer to that question is no.²⁰ What justice requires is not in part a function of what people can bring themselves to do. More generally, on his view it is a mistake to think that a “hopeless” theory of justice, which we have no good reason to think will ever be implemented, even though it could be, is, in virtue of being hopeless, worse than a “hopeful” theory of justice, that is, one that we have no strong reason to think will not be successfully implemented.²¹

Though the idealists reject the standard of practicality, they seem to agree with the implementers in key respects on what it means for a theory to be practical, that is, the IP conception of practicality. So in his discussion of fundamental principles of justice versus feasible rules of regulation, Cohen says that, unlike for rules of regulation, “it is not a constraint on a sound conception of justice that it should always be sensible to strive to implement it.”²² Rawls was not “investigating the nature of justice as such” but *instead* trying to “establish what the right principles to implement socially are,”²³ i.e., rules of regulation. We can see here Cohen agreeing on part (a) of IP, that a practical theory is ultimately implementable. Part (b) of IP says that a practical theory is one that provides a sufficiently accurate picture of what will happen when the theory is imple-

¹⁷ G. A. Cohen, *Rescuing Justice and Equality* (Cambridge, MA: Harvard University Press, 2008), 84.

¹⁸ *Ibid.*, 86.

¹⁹ David Estlund, “Human Nature and the Limits (if any) of Political Philosophy,” *Philosophy and Public Affairs* 39, no. 3 (2011): 207–37, at 226.

²⁰ “The content of justice . . . might . . . precede any facts about what humans or any beings happen to be like” (Estlund, “Human Nature,” 229), including facts about their altruistic tendencies or abilities to comply with particular laws or support particular institutions.

²¹ David Estlund, *Democratic Authority* (Princeton, NJ: Princeton University Press, 2008), 264–67.

²² Cohen, *Rescuing Justice*, 271.

²³ *Ibid.*, 301.

mented. To know this is to know certain facts about society, facts that, while on Cohen's view are irrelevant to the substantive content of justice at a fundamental level,²⁴ are nonetheless crucial for determining the practical question of which rules of regulation we should adopt.²⁵ Similarly, Estlund makes a distinction between what justice requires and "what we should do, in practice, given what people and institutions are likely to do."²⁶ Similarly to Cohen, he adds, "we need to concede the facts in practice, even if not in our moral conclusions."²⁷ When it comes to justifying our choices regarding the practical matter of what to do, we will need to take into account the relevant facts. If we do not, we may try to do things that are futile, wasteful, or harmful,²⁸ that is, things that oppose or detract from successful implementation. But when it comes to identifying what, in principle, justice requires, we need not concede anything to the facts.

At one point, Estlund seemingly resists the idea that if a theory is "hopeless," it follows that it is not practical. A theory could be hopeless because an insufficient number of people will bring themselves to comply with it, even though they could. Such a theory "counsels all people, together, to behave differently"²⁹ and would count as practical because it "does prescribe action in a certain way."³⁰ On the face of it, it may seem as if Estlund is offering an alternative account of practicality, thus undermining my claim that idealists, though rejecting practicality, agree with implementers about what practicality consists in. To allay this concern all we need to do is recall the distinction, set out in Section I, between formal and substantive practicality. What Estlund shows here is that hopeless theories can be formally practical: they can tell us to act. But formal practicality is not an alternative to conceptions of substantive practicality like IP, which describe the relationship that the prescribed acts are supposed to have with the world in order for the acts to be successful. It does not appear that Estlund provides us with an alternative account of that, in which case we have no reason to think he rejects IP as the appropriate conception of practicality.

Cohen says, "the question for political philosophy is not what we should do but what we should think, even when what we should think makes no practical difference."³¹ That may be pithy enough for an idealist slogan, but just in case we need something shorter with which the idealists can reply to the implementers, it could be this: practicality is not justice.

²⁴ Fundamental principles of justice, according to Cohen, are "in no way dependent on the character of any facts." *Ibid.*, 281.

²⁵ *Ibid.*, 272.

²⁶ Estlund, *Democratic Authority*, 268.

²⁷ *Ibid.*

²⁸ *Ibid.*, 266.

²⁹ *Ibid.*

³⁰ *Ibid.*

³¹ Cohen, *Rescuing Justice*, 268.

IV. THE IMPLEMENTATION-PREDICTION CONCEPTION AND EPISTEMIC OVERCONFIDENCE

In the previous two sections I have argued that both the implementers and the idealists agree on what it means for a theory of justice to be substantively practical. More specifically, both agree with

IP, or the *implementation-prediction* conception of *practicality*:

A principle or set of principles is “practical” if (a) its implementation by way of a series of events that begins at the relevant status quo and ends at something relevantly similar to what the theory prescribes is both plausible and desirable, and (b) the predictions (both descriptive and evaluative) about its implementation which contribute to its justification are sufficiently accurate.

The idealists reject the idea that theories of justice have to be substantively practical. The implementers, of course, disagree. In this section of the essay I will argue that both the idealists and the implementers are mistaken in their respective stances towards practicality. Their mistakes are rooted in the same problem: thinking that we know more than we do.

A. Why the idealists should not reject practicality

The idealists seem to be committed to the following claims: (1) we can know whether principles of justice are correct, and (2) principles of justice need not be practical.³² (Insofar as the idealists’ substantive conception of practicality is IP, then a principle of justice that is not practical is a principle that may never be, or may never have been, implemented. And if a principle is never implemented, we never have experience with it.) It follows from these that (3) we do not need to have experience with the correct principles of justice in order to know that they are correct.³³ I think we have strong reasons to reject claim (3). If we reject (3), then either claim (1) is false, claim (2) is false, or both are false.

In arguing against (3), I am committed to the view that in order to know whether a principle of justice is correct, we have to have experience with it. This is an epistemological point; I am not taking any position on the metaphysical status of correct principles of justice. Furthermore, I am not defending a full-blown empiricism here. It may be possible that some things are knowable without having some related experience, but principles of justice are not such things.

To see why, it might be useful to start with something simpler than a principle of justice. Consider, instead, the following evaluative claim, W1:

³² As a reminder, unless otherwise noted, “practical” refers to “substantively practical.”

³³ If you object to using the term “correct” to describe principles of justice, substitute “most acceptable” for it.

“this wine is excellent.” There are some standards for excellence in wine, and W1 asserts that the wine in question meets these standards. No one would expect that the truth of such a claim could be made apart from experiencing (seeing, smelling, and drinking) the wine. Now move up a level to the claim W2: “Excellence in wine consists in X, Y, and Z” where X, Y, and Z are three specifications of important factors in assessing the quality of a wine. W2 is a claim about what the standards of excellence in a wine are. Can we know the truth of W2 without ourselves (or someone whose testimony we are relying on) having had experience of wine? We might be able get started without experience. We know wine is something we taste, and so excellent wine will at least taste good. Since it is something we put into our bodies, it should not be toxic (or too toxic). What else? It is a liquid, so it should pour. That does not take us too far down the path toward excellence.

What are the kinds of things critics and those who drink a lot of wine say about excellent wines? Among other things, they have aromas and tastes that are complex (multiple flavors together), multilayered (different flavors distinctly emerge as the wine makes contact with different parts of the mouth), interesting (possessing not just the basic multiple flavors, but unusual ones), balanced (the ratio of fruit flavors to acidity is such that the wine could develop further, in a good way, if aged further in the bottle), and, of course, delicious. Suppose the critics are right that these are valuable features of a wine. How did they come to determine this? More to the point, could they have come to identify these valuable factors apart from any experience with wine? No.

Importantly, we need experience with wine to know what wine is capable of, what it can be like. And we need to know what wine can be like in order to know what standards it is appropriate to hold it to. That wine can be complex or multilayered—and to what extent—is not knowable *a priori*. Our sense of what deliciousness for a wine can be is informed by our understanding of these other factors, as well as the range of flavors and aromas and textures that it is possible for a wine to have. That is, there is a dynamism at work. We discover, perhaps inadvertently, certain aspects of wine, such as complexity, and we then identify complexity as a part of the evaluative standard for wine. In turn, we seek out complexity, come to appreciate it, enjoy it, find it adds to deliciousness, see that it makes multilayeredness possible, and we come to appreciate that, too. And so on.

We learn about and incorporate these different factors into our evaluations. But how do we organize these factors into an evaluative standard? It is unclear how, apart from trying different kinds of wines, one would be able to assess the relative importance of the different factors in determining which wines are excellent. Even if there are some facts of the matter regarding what makes a wine excellent, our understanding of what it means for wine to be excellent is rooted in our experience with

wine. We might posit what we take to be the appropriate standard of excellence, only to realize, say, through encounters with new wines, or new foods or beverages, or new ways of drinking or storing wine, that it suffers from false positives and false negatives. We then revise our standards in light of this information.³⁴ In short, we use empirical information to both formulate and confirm our evaluative standards.

When we say that a particular society is just, we could be making a claim that is analogous to W1, the claim that some wine is excellent. That is, we could be presuming some criteria for the justness of a society, and asserting that the society in question meets the criteria. As with W1, no one would doubt that we must provide empirical evidence in substantiating this claim.

When we say what the standards for a just society are—that is, when we articulate the principles that a society must conform to in order to be just—we are making a claim that is analogous to W2, the claim about what excellence in wine consists in. And as with W2, it is difficult to see how we could make this claim absent some experience with the relevant phenomena the principle is about, and with the principle itself.

To see this, suppose we determine that one member of the set of correct principles of justice requires, in some as-of-yet unspecified way, equality. For this principle to be at all informative, let alone of any use, it must specify the respect in which (or the space within which) people should be equal.³⁵ Equality before the law? Equality of resources, and if so, which ones? Equality of liberties? Income? Welfare? Capabilities? Opportunity? Political voice? Equality for all individuals, or just citizens? Cultures? Religions? Sexes? How do we learn which spaces are ones that demand equality, and for what kinds of social entities? If we cannot have all equality in all of these ways at once—which we of course cannot—what are the relative weights of the kinds of equality on offer? How can we figure that out so we can make wise trade-offs when necessary? (Similar questions come up for other basic political values besides equality, of course.)

As with wine, there may be some general, vague, claims we can make about the kind of equality justice requires without having any particular experience living with a principle of equality. Recall, though, that with

³⁴ Again, this point is neutral between realist and constructivist accounts of evaluative standards. We could be revising our standards in an attempt to achieve mere reflective equilibrium, or we could think that it is only through investigation that we have a chance of properly identifying the mind-independently true evaluative standards. If we are constructivists of some sort, then of course we would need to try them out to agree on them in any informed way. But if we are realists, *that the principles seemed right to us when we tried them* would merely be good evidence that we have discovered the right principles. Given that we are often wrong in our predictions about what works, we should not trust our armchair judgments.

³⁵ Amartya Sen, *Inequality Reexamined* (Cambridge, MA: Harvard University Press 1992) chap. 1.

wine, to have any confidence in *specific* evaluative claims, someone needs to drink a lot. Likewise, to get a more specific sense of what justice requires vis-à-vis equality, we need to try out different kinds of equality. By doing so, we will learn what it is like to be treated equally in those different ways. We might learn about new ways of treating people equally, or about important inequalities that develop when aiming for certain kinds of equality. We may come to recognize certain spaces for equality as more important than others. Our judgments about equalities or inequalities may change as a result of experiencing them. It may be that inequalities we initially object to are not bad in the ways our theory predicted, or that equality of certain kinds is not desirable, or that some inequalities are more important than others for showing people respect. Perhaps by living under certain conditions of equality we will come to have a new appreciation for its value, and be able to offer better justifications for it, or perhaps we will come to reject equality in most forms; in either case, our fundamental criteria of justice may alter, and our view of what makes for a desirable outcome may alter, too.³⁶

Political philosophers—idealists and implementers alike—have a view of what life is like in the absence of their principles, and a view about what it would be like if their principles were in place. If they learned that these views were mistaken, that could affect their specification of, or support for, their principles. Additionally, these views about what the world would be like (namely, how good or bad it is, or whether it measures up to the correct normative ideals) can be informed by experience with the principles the views are supposed to support, which can lead to the refinement and specification of the principles. Further, through this learning experience, our criteria for how we go about identifying a good principle of justice may themselves change—they are not fixed. Because so many variables relevant to our understanding of justice can be affected by our experience, it seems implausible that we could determine which principle or set of principles aligns these variables correctly without having experienced them.

It may be objected that there are certainly *some* normative claims we can know to be true without experience, for example, that slavery is wrong. So firm are our convictions on such matters that it may seem that they could not depend on something as contingent as our experience. While a full defense of the quasi-pragmatist approach to the epistemology of justice is beyond the scope of this paper, I will say a few things in response that may address some common concerns. First, let me reiterate that the experiential view advanced in this section is about how we come to *know* certain normative claims, and not about what

³⁶ My view on these matters is influenced by Elijah Millgram, *Practical Induction* (Cambridge, MA: Harvard University Press, 1997) and Henry Richardson, *Practical Reasoning about Final Ends* (Cambridge: Cambridge University Press, 1994).

makes these normative claims *true*. My view is that the epistemology of justice is empirical, but that does not imply that the metaphysics of justice is contingent. So if one's concern is that our set of bad experiences with slavery is too flimsy or fickle a basis for an *explanation of the wrongness* of slavery, that is compatible with my view, which holds only that those experiences are a necessary part of the *explanation for how we came to know* that slavery is wrong. Second, the strength of our current judgment that slavery is wrong could be explained by our experience of its awfulness, and the awfulness of its effects, over a long time. We have seen the damage that slavery does not just to those enslaved, but to cultures, damage that persists generations after its cessation, and we have had more time to absorb the lessons of this experience. Third, this experience-based explanation for the strength of our opposition is consistent with the common view that early wrongdoers are less deserving of opprobrium. Compare the moral disapproval one feels towards a slaveholder in ancient Greece, a slaveholder in the antebellum United States south, and a slaveholder in present-day Manhattan. We shrug our shoulders at the first, shake our heads with disapproval at the second, and explode with outrage at the third. At this point in history we know better, and we have no excuse. Fourth, personal experience is not the only source of moral information. If we have determined, through experience, that some act or practice is wrong, we can conclude that whatever other acts and practices are entailed by the ones we know to be wrong are themselves wrong, and further, that other acts and practices that are relevantly similar are also wrong. Additionally, the relevantly similar experiences of relevantly similar others may be made use of, as well. (Of course, what counts as "relevantly similar" may itself be something we come to grasp with the assistance of experience.)

Much more can be said to defend this role for experience. If we accept it, we have good reason to reject claim (3), the idea that we do not need to have experience with the correct principles of justice in order to know that they are correct. Claim (3) followed from (1) we can know that some principles of justice are correct, and (2) correct principles of justice need not be practical. Therefore, we must reject (1) or (2).

Owing to the ways in which experience can inform and overturn our evaluative judgments, and to the fact that potential experience is limitless, I think we have a more tentative view of what we know about principles of justice than is expressed by (1). Our understanding of our knowledge of principles of justice should be based on the model of scientific knowledge, and that means less certainty, not more. For any accepted scientific claim, the claim explains the relevant data, fits with other scientific claims that are likewise explanatory, and is held firmly but tentatively owing to the possibility that new, uncooperative data might arise. It is unclear why ethicists and political philosophers would be entitled to greater certainty

in their claims about the world, than, say, physicists are. But this is less a rejection of (1) than a reinterpretation of what it means to know that a principle of justice is correct.

If we reject (3) but accept (1), then (2) must go. Correct principles of justice have to be practical. Of course, this is exactly what the idealist rejects. The central idealist point is that principles of justice may be correct, even if we lack the ability or will to put them into practice. My argument may seem like a rejection of the idealist position, then. It is not. As we will see in Section V, there is a conception of practicality that is compatible both with the rejection of (3) argued for in this subsection, *and* the acceptance of the central idealist point that, for principles of justice, implementability does not constrain truth.

B. Why the implementers should reject the implementation-prediction conception

While the idealists are mistaken in rejecting the idea that principles of justice should not be held to a standard of substantive practicality in general, the implementers, I will argue, are mistaken in accepting the IP standard of practicality. What leads idealists to make their mistake is a kind of epistemological overconfidence, the idea that we can know what principles of justice are correct without having some experience with them. What leads the implementers to make their mistake, I will argue, is also a kind of epistemological overconfidence, specifically, an overconfidence in our abilities to determine what is possible, and to identify what we will come to judge as desirable. We are justified in having a fair amount of skepticism regarding these abilities.

Let us start with our capability for determining what is *possible*.³⁷ The first two things to consider here are history, which is full of technical and cultural predictions that have failed to come to pass, and the present, which is full of technical and cultural developments that no one would reasonably be thought to have been able to predict. I would go further and venture that for most technological or cultural developments today, there were many populations at some earlier point in time about which we could plausibly say they had no reason to think these things could occur. It is a reasonable induction to suggest from our history of mistaken, incomplete predictions that today's predictions are also mistaken or incomplete. Because *we are our ancestors' seemingly impossible future*, we have good reason to think *we are ancestors to a seemingly impossible future*.

According to the implementers, a constraint on the acceptability of a principle of justice is practicality, and practicality is understood as requiring implementability. So, to know whether a principle of justice is correct we need to know whether it is possible to implement the principle. How-

³⁷ Discussed in Weinberg, "A Little Reality is a Dangerous Thing."

ever, we do not know what is possible, especially given a long enough time frame. So we cannot know whether the principle is implementable, and so we cannot know whether it is correct. In such cases this may render the IP standard itself impractical. As such, the implementers should reject it.³⁸

Now let us turn to our capacity to determine whether the implementation of a principle of justice is *desirable*. Here, too, our capacity is hampered by our limited ability to predict what will happen. We do not know what life under a particular principle of justice would be like, or what the execution of the principle demands in terms of our political institutions, or more generally what consequences would follow over time from having the principle in place. Since our view about the desirability of a principle might be affected by these things, our ignorance of them—call it *descriptive ignorance*—contributes to our uncertainty.

Descriptive ignorance is only one factor in our uncertainty. Another factor is that we are creatures who change our minds; more specifically, we alter our evaluative judgments in light of our decisions and experiences, sometimes in unpredictable ways. So even if we assumed we knew what was going to happen in the world, we do not know what we are going to think about it. Even if we have a prospective inkling of what we will come to think of something, the experience of that thing may be transformative. Afterward, we may come to have a completely different view. Consider the person who does not like opera but begins going often in order to please a relative. We can imagine that as she learns more about opera and how to appreciate it, she enjoys it more, and comes to have a different evaluation of it than she ever expected. Sometimes we just do not know what our judgments about something new will be; call this *evaluative ignorance*.

Evaluative ignorance scales up from individuals to societies. Living in a society governed in accordance with egalitarian principles may, over the course of several generations, have an effect on what people think about egalitarianism. Or consider trade-offs between liberty and security. We can envision sets of principles that rank these goods differently. One set of principles leads to the “liberty” society, the other to the “security” society. After the institutions of these societies have been up and running, and running well, for a while, do we think that the inhabitants of the liberty and security societies will answer surveys about their political values in the same way? No. Their attitudes will be transformed by living under different principles and institutions.

The changes to our evaluative attitudes may not be predictable. Again, we can look at examples of past novelties to see how wrong we some-

³⁸ Note that changing the language from “impossible” to “unlikely” does not help the implementers, as likelihood is going to be prospectively judged, based on an interpretation of the facts—an interpretation we have reason to distrust.

times were about what we would think were something to come to pass. Compare “what people think they will think once X happens” with “what people think, now that X has happened and they have lived with it for a while” for a number of Xs: women’s suffrage, slave emancipation, in vitro fertilization, mobile phones, and antidepressants, for example. In each of these, evaluative attitudes that were common prospectively became rare or rejected retrospectively.

Taken together, descriptive ignorance and evaluative ignorance present substantial obstacles to putting forth only those principles both that we prospectively know can be implemented and that will have the kind of effects that we will then come to approve of. We are just not in such a privileged epistemic position. So the IP-standard, which is supposed to check unrealistic political philosophizing, itself is unrealistic.

By rejecting practicality, the idealists presume that we can know, absent relevant experience, that a principle of justice could be correct. By accepting the IP standard of practicality, the implementers presume that we can know, absent relevant experience, that a principle of justice cannot successfully be implemented. As we have seen, we have reasons to reject both presumptions. Contra the idealists, we should accept practicality, but, contra the implementers, we should not accept the implementation-prediction conception of it.

V. THE EXPERIMENTATION-LEARNING CONCEPTION OF PRACTICALITY

The conception of substantive practicality we should make use of is what I call the *experimentation-learning* conception, or EL. It can be built into a standard we then apply to principles of justice, a standard that both idealists and implementers have reasons to accept, and that generates specific roles for both ideal and nonideal theory.

According to EL, a principle is “practical” if (c) experimentation with the principle in order to determine whether it is plausibly implementable is currently possible, or foreseeably possible, and (d) it is reasonable to think it is possible to learn from said experimentation whether the principle is one we would endorse once we had the experience of living with it.

A crucial difference between the IP and EL conceptions of practicality is epistemological. Where IP speaks of implementation (a), EL does not presume that implementation is possible and speaks instead of experimentation (c). Where IP speaks of correct prediction (b), EL speaks of learning (d).

When built into a standard for assessing principles of justice, IP asks that we know a lot more about what could be called the “success” of a principle prior to trying it out. The EL standard, in contrast, only asks whether it is possible to experiment with, or try out, a principle and then learn from the experience. That is much less epistemologically demanding.

A rejection of a standard of practicality altogether implies that we could know the correctness of principles of justice without ever having to try it. This assessment of our evaluative and predictive capacities is overoptimistic. Rather than assume we already have this information, or could know about the acceptability of a principle without it, holding principles to the EL standard is meant to elicit this information.

The EL standard is thus more sensitive to the epistemological challenges of assessing principles of justice. It does not ask us to know more than we do, and it gives us a means by which to come to know what we need to know.

While the argument for EL shows that both the idealists and the implementers make a mistake, EL also allows them both to be partially correct. The idealists are rightly concerned that a standard that demands implementability is inherently conservative and is likely to rule out more just alternatives because they differ too much from the status quo. The EL standard does not demand implementability, only experimentation. Further, in demanding experimentation, The EL standard asks only that it be plausible to try out a candidate principle of justice, not that the results of the experimentation be positive; so experimentation is not a version of successful implementation. Meanwhile, the implementers are right to think that information about how a principle would fare in the real world is relevant to our knowing whether we should accept it, and the experimentation that the EL standard demands is aimed at providing that information.

In Section I we saw that ideal theory is criticized by implementers as being impractical because it relies on unrealistic assumptions that will yield principles that cannot be successfully implemented. This criticism has force only if principles generated in ideal theory are to be implemented. The idealists who make use of ideal theory have replied by denying that principles of justice need to be practical, and so, denying that they need to be implementable. This may leave us wondering what the products of ideal theory are for. What are we to do with ideal theory? EL provides an answer.

According to EL, principles of justice generated under ideal theory are candidates for experimentation. They are *for* trying out. We need not assume that they can be successfully implemented, nor that they will lead to the outcomes ideal theory predicts, nor that the outcomes they lead to—predicted or not—will be ones we would now or afterward come to judge as good or just. Ideal theory, I would argue, is a particularly good approach to developing candidate principles with which to experiment.

Given that philosophers have only been considering principles of justice for not quite two thousand five hundred years, and that humanity and successor forms of intelligent life could exist indefinitely into the future, it would be quite a surprise if, in this early stage of history, we had already come across the best principles of justice.³⁹ The “unrealistic” assumptions of ideal theory allow us to better resist various biases and heuristics that limit our ability to imagine alternatives that differ substantially from the status quo or previously introduced alternatives.

EL allows for the proposals of political philosophy devised under non-ideal theory to have a similar role as candidates for experimentation, though EL makes clear that nonideal theory does not have exclusive bragging rights to practicality—at least the kind of practicality that we should be interested in. Nonideal theory’s proposals may be closer to the status quo than those of ideal theory, but that in itself does not give us a reason to think they are more likely to be true. EL also provides a new role for nonideal theory, or perhaps points to a need for a new kind of non-ideal theory. The task of this nonideal theory would be to help determine which principles are testable, or to formulate versions of them that are testable, and to determine methods of learning from experimentation with them. In other words, nonideal theory’s focus can be directed away from the question of which principles are realistically successfully implementable, to the question of which principles are realistic candidates for experimentation.

VI. OBJECTIONS TO AND EXPLANATIONS OF EXPERIMENTATION

The idea of experimentation put forth in the EL conception of practicality needs further development, and there are clearly a number of questions about how the experimentation itself is to be implemented. While I leave the development of a full-fledged account of the experimentation to another time, I will mention some objections here and sketch answers to them. In doing so I hope to provide a somewhat fuller picture of the kind of experimentation I think is worthwhile.

One objection is that the experimentation idea is unrealistic. How exactly could we experiment with principles of justice?

The first point I want to make in response to this question is to emphasize that my argument does not, strictly speaking, direct us to engage in experimentation. Rather, the point of the argument has been to show that the dispute over holding theories of justice to some test of practicality has thus far been misconceived. Insofar as “practical” has meant something

³⁹ Derek Parfit famously says of nonreligious ethics, that, because it is so young a science, “it is not irrational to have high hopes” for its success. Parfit, *Reasons and Persons* (New York: Oxford University Press, 1984), 453–54. I add here that because it is so young, it might be unreasonable to think that it has already succeeded.

like “implementable,” it turns out that both those who favor and those who oppose practicality have been holding theories of justice up to epistemologically inaccessible goals, regardless of whether one’s preferred approach is ideal or nonideal theory. Political philosophers would do better to understand their theories as proposals for us to experiment with and learn from. But these points do not add up to an imperative to experiment.

That said, insofar as we recognize the epistemic position we are in regarding justice, and insofar as we care about justice, it may behoove us to try out some principles of justice as experiments to learn from. And since societies are not petri dishes in a lab, there are some important questions about this. Indeed, there are a variety of political, social, and economic obstacles such experimentation would face. Helpfully, we have a model that might guide us as we try to envision and work out the problems of experimentation: federalism. As Justice Brandeis famously wrote, “It is one of the happy incidents of the federal system that a single courageous state may, if its citizens choose, serve as a laboratory; and try novel social and economic experiments without risk to the rest of the country.”⁴⁰ The idea of the states as laboratories does not accurately describe contemporary U.S. society, and to revert to a more experimental federalism would face many challenges. Even if U.S. federalism were to be more experimental, it may not be sufficiently so, given the Constitution’s status as the supreme law of the land. Still, radical federalism may be one way of executing experiments in justice. And we should note that many of the obstacles involved in such a proposal are no greater than those faced by any normative political theory that sets out an ideal polity that differs from our current one, as most theories do.

A second objection is that the experiments will be impossible because of a lack of guinea pigs. We would not want the experiments to involve forcing persons to live in certain places. So who would want to move into, or continue to live in, one of the less successful “experiments”? First, though I have emphasized the value of experimentation, we should not forget that more ordinary modes of critique remain available to us, as I mentioned earlier. We can check for the internal coherence of sets of principles, for example. We can rule out any arrangement previous experience with which has shown to be morally unacceptable. And we will probably use our current considered judgments to set some limits on the degree of experimentation. Still, some experiments will turn out better than others, and inhabitants may indeed wish to move from, or end, an experiment. If enough so desire, that particular experiment may cease. This is not a problem from the point of view of the EL standard. One particular datum the experimentation helps supply is what it is like to live in a certain kind of society. If the persons who are living there find

⁴⁰ Louis J. Brandeis. *New State Ice Co. v. Liebmann* (dissenting), 285 U.S. 311 (1932).

themselves unable or unwilling to continue, that is useful information. As in any experimental setting, failure is a result from which we can learn.

A variation on this second objection is that, in creating so many different kinds of societies, experimentation risks harming persons, especially since in our experiments we will likely more often fail than succeed at achieving acceptably just societies. It is indeed true that the experimentation risks harming persons, if we presume that injustice correlates somewhat with increased harm. But the question to ask, then, is whether our current society is just. If we are fairly confident that it is, then the need for experimentation is slight, and there will be few experiments, or the experiments will not deviate wildly from the status quo; so the risk of harm is not great. But if our current society is not just, then the failure to engage in experimentation will lead to persons being harmed—in our society—for the failure to experiment is likely to delay our recognition and achievement of justice.

A third worry is that experimentation *uses* people, treating them as mere sources of information about justice, institutions, and intuitions, rather than as autonomous beings with dignity. In response, note that it is likely that under the kind of experimentation being suggested, persons would have a choice about which experiments to take part in—after all, if a principle of justice is to gain our allegiance, it must resonate not only with our understanding of justice once we are fully immersed in its resultant culture, but it must be attractive or intriguing in some way for some of us before the experiment gets started. Otherwise the principle would never be a live option for us to even strive toward. Thus, insofar as people get to choose whether they are to live in an experimental setting, or choose which to live in, they are treated as beings with their own ends.

There are additional worries related to the capacity of such experimentations to serve as actual experiments from which we can learn. Some of these are methodological concerns that I am not equipped to address, or are best answered in the context of particular experiments, and relate to the challenges involved in data collection, isolating the relevant variables, identifying the right causal relations, avoiding distorting factors, generalizing from the results, and so on. To take seriously the idea that principles of ideal theory could serve as the basis for social experimentation would involve a commitment to doing our best to address these methodological issues.

A broader concern that someone might raise about experimentation is why we should give more weight to the opinions of persons living under the principles of justice in question in an experiment, than we do to the opinions we currently have, prior to or outside the experiment. After all, the question we would be using experimentation to answer is what principles of justice *we* should adopt. Suppose our views about justice might, as a result of living in one of the experiments, come to be different than

they are now, and that we are strongly committed to our current views. Doesn't our attachment to our current views give us a reason not to experiment? But this takes us back to the very reasons for rejecting IP and accepting EL. There are things we can only learn about a principle, and about a society, by experiencing them. The experiments do not merely *change* one's opinion about conceptions and institutions of justice, if they do; rather, by giving one a broader range of experience on which to judge, they *improve* one's opinion. However, I admit that clearly more needs to be said here, particularly in light of worries about adaptive preference formation.

Finally, there is a question about whether we can ever actually experiment with principles of justice. There are different reasons to think the answer may be that we cannot. One reason may be that, because principles of justice make contact with the world via policies, what we are really experimenting with are policies, not principles.⁴¹ The same principle can be instantiated by different policies. Suppose our experiment yields negative results. It could be argued that this should count only against the policies in place, rather than the underlying principles of justice, since those policies could be swapped out for others that also instantiate the principle. To put it another way, there is a question of whether what failed was the theory, or just the way it happened to be put into practice.

This objection shows that experimentation may be a much more drawn out process than one anticipated. It may be that we have to test multiple policy instantiations of principles of justice in order to be warranted in drawing conclusions about the principles themselves. That there are countless possible policy instantiations of any given principle is not an insuperable problem, though. Consider scientific experimentation. Suppose we are testing a hypothesis about the interaction of light and metals to see whether it generally (if not exceptionlessly) holds. Since there are many different kinds of light and many different kinds of metals, it may be insufficient to attempt to confirm or falsify the hypothesis with just one experiment with one kind of light and one kind of metal. But it would also be false to claim that the only way of successfully testing the hypothesis is by testing every possible combination of light and metal. There will be a sample of sufficient size and diversity from which it is reasonable to draw conclusions about the hypothesis. Likewise, we should conclude that there could be a sufficiently large and diverse set of policies that instantiate a principle of justice, with which we could experiment, and from which it would be reasonable to draw conclusions about the principle. One experiment would perhaps be too few, but to think that only an infinite number of experiments is sufficient would be to hold experimentation with justice to a much higher evidentiary standard than we employ in any other inquiry.

⁴¹ This version of the objection was put to me by Chris Bertram.

VII. CONCLUSION

Must principles of justice be practical? That question is typically understood to be asking about IP, that is, about whether we could successfully implement the principles of justice and, in doing so, bring about the predicted just society. I have argued that this is a faulty way of understanding practicality. The surefooted implementers who tend to endorse practicality cannot in fact meet the IP standard, while the starry-eyed idealists who tend to reject practicality in general cannot do so on the basis of problems that are specific to IP. While we should indeed reject the view that principles of justice must meet the IP standard of practicality, we should accept a standard that holds them to the experimentation-learning conception of practicality. Understanding principles of justice as hypotheses to try out and learn from better reflects what we do know and what we need to know so that our political philosophy can reasonably aim at what is both practical and true. In other words, if we are wondering whether, and to what extent, political philosophy has room in it for the real world, we can find out by trying to make room in the real world for political philosophy.

Philosophy, University of South Carolina