

Why Political Philosophy Should Be Robust

ALEXANDER S. KIRSHNER *Duke University, United States*

JEFF SPINNER-HALEV *University of North Carolina at Chapel Hill, United States*

Political philosophers and theorists make arguments about high-stakes problems. This article shows that those theories would be more credible if political philosophers ensured their work was robust: capable of withstanding reasonable changes to their assumptions and to the cases to which their arguments apply. The world is varied and inconstant. As a result, scientists and social scientists recognize the virtue of robustness. This article shows why political philosophers should also do so. It defines robustness, demonstrates its value, and shows how it can be evaluated. Illustrating the stakes of robustness, the article assesses prominent arguments concerning multiculturalism and open borders. Avoiding misunderstanding and confusion should be a central aim of political philosophy. To sidestep these outcomes and to reassure scholars that one's theory is not subject to concerns about its credibility, it will often be reasonable for philosophers to explicitly test their theories for robustness.

INTRODUCTION


Mass migration, climate change, racial injustice, extremism, and inequality are defining features of this moment. And political theorists and philosophers have energetically tackled these issues, seeking to generate warranted conclusions about the most pressing problems that societies face (Anderson 1999; Caney 2012; Carens 2013; Fricker 2007; Müller 2016; Shelby 2016; Urbinati 2014; Young 2011). They do so by trying to satisfy the virtues of sound argumentation: forging theories whose conclusions both follow directly from their premises and do not violate other core moral and political principles. That is no small task. But in this article, we spell out a distinct theoretical virtue, a characteristic of a theory providing scholars a reason, perhaps not dispositive, to prefer that theory over another (Kuhn 2013). And we think this virtue should be of concern to all political theorists. It is robustness.¹


A theory is robust if it remains credible when considering the range of cases and circumstances targeted by the author, as well as other philosophically similar cases—that is, cases and circumstances to which we would reasonably expect the theory to apply. By

credible, we mean the degree to which a theory warrants belief by virtue of the evidence and arguments it marshals.²

On the basis of our definition, robust theories might target a very broad range of circumstances (e.g., arguments for the general superiority of democracy), as they often do, or a very narrow range of circumstances (e.g., arguments outlining a specific solution for a particular circumstance). Our focus is on the former. But there is nothing inherently amiss with the latter. Scholars are often explicit about the cases and circumstances they are targeting—for example, a theory might apply to all democracies or just democracies with significant ethnic minority groups. Other times, they are not explicit about this and readers must use the theory itself and the context to determine which cases and circumstances are under discussion. Regardless of a theory's breadth, at a minimum, robustness means a theory remains credible when tested across the range of situations an author is talking about.

Why should robustness matter to political philosophers? Political philosophers often tackle big topics. And they aim to shed light on a world of profound complexity, variation, and inconstancy. To gain traction on their topics, they often consider their subjects under a simplified and restricted set of conditions. Proceeding in this manner is entirely reasonable and familiar enough. Yet this approach carries a persistent risk. The ostensible credibility of a theory that explores a broad subject under particular, restricted conditions may actually depend, problematically, on the particularities of those conditions. That is, the theory might be problematically ad hoc (Douglas 2013, 800–1). Another way to describe this is as follows: political philosophers

Alexander S. Kirshner , Associate Professor, Department of Political Science, Duke University, United States, alexander.kirshner@duke.edu.

Corresponding author: Jeff Spinner-Halev , Kenan Eminent Professor of Political Science, Department of Political Science, University of North Carolina at Chapel Hill, United States, spinner@email.unc.edu.

Received: April 10, 2022; revised: January 06, 2023; accepted: August 08, 2023. First published online: October 16, 2023.

¹ Throughout this essay, we use the terms—political theorists and political philosophers—interchangeably. In our view, robustness should be considered a virtue no matter which group or groups one identifies with.

² We focus on credibility rather than truth because theories that are poorly constructed and not credible might be true for reasons the theory does not comprehend. And political theories can be credible even if they ultimately prove to be false.

understandably use a simplified set of premises when developing arguments meant to shed light on complicated real-world conditions and cases. Theories lacking robustness are not fully credible when the premises of the argument are shifted to better match the real-world conditions and cases an author is targeting.

What is the price of unrobust theories? Because such theories can appear to be robust across a range of cases and circumstances, they can lead both authors and readers astray, causing them to significantly overestimate a theory's import and misinterpret its implications, to spend time and effort defending or even building on claims that should not be depended on, like erecting a seafront condominium on a cracked foundation. Since the goal of political philosophy is, presumably, to generate insight, avoiding misunderstanding and confusion should be a central aim of political philosophers. To sidestep these outcomes and to reassure scholars that one's theory is not subject to these concerns, it will often be reasonable for philosophers to explicitly test their theories for robustness. We spell out the logic of two such tests: an *alternative assumptions test* and an *alternative cases test*. We think employing these tests will produce better political philosophy, political philosophy that is more credible and more likely to generate insight into how to tackle the challenges and opportunities of the political moment.

Scientists and social scientists regularly work to identify robust models and explanations of phenomena, models and explanations that are not undermined when alternative methods, assumptions, and cases are considered. Systematically testing the robustness of one's work is critical for assuring other scholars that the models and explanations outlined are credible and that ostensible conclusions are not ad hoc artifacts of a particular model specification or a particularly advantageous case. When models and explanations are shown to be robust, readers of that work gain assurance that the research effectively captures something about the targeted phenomena. For these important reasons, the *American Political Science Review* requires authors to rigorously test their work using alternative models, scope conditions and measures, amongst other steps, because "empirical APSR papers are expected to be robust" (Editors 2022, vi). This article argues that similar expectations should apply to political philosophy. The reasons to embrace robustness as a virtue in the sciences and the social sciences—the import and the difficulty of making credible claims about a complex, varied, and changeable world—apply widely, if not universally, in the context of political philosophy.

In the next two sections, we fill out our definition of robustness, identify why it is a virtue of political philosophy and outline two ways of testing for it. We then demonstrate the clear practical utility of our tests by examining two powerful theories. The first, advanced by Will Kymlicka, focuses on the justification of state support for minority national cultures. The next theory is outlined by Michael Blake. He seeks to show that citizens of a polity may exclude immigrants to avoid the imposition of novel moral

duties. We illustrate how these theories' robustness can be tested. Our analyses demonstrate why an explicit assessment of robustness can lead to a better appreciation of these theories' implications and better political philosophy.

ROBUSTNESS AS A VIRTUE OF POLITICAL PHILOSOPHY

Whether one is discussing a scientific measurement or the stability of an apartment tower, robustness means that something remains unchanged despite the fact that some factor, for instance, a new measuring device or an earthquake, could have changed it (Calcott 2011, 289; Levins 1966; Wimsatt 2007). A *marriage* is robust if it persists after partners experience changes in their economic fortunes and personal desires. A *conclusion* about the world is robust if several different and independent studies using different data and distinct methodologies arrive at a common finding. A *statistical model* is robust if estimated effects remain reasonably stable when regressors are added or when alternative, plausible model parameters are tested.³

Here, we focus on the robustness of political theories. A *political theory* is robust if it remains credible when considering the range of cases and circumstances targeted by the author, as well as other philosophically similar cases. It is plain to see why we would expect a theory to retain its credibility when tested against the range of cases and circumstances targeted by an author. But there is a second, perhaps less central set of cases and circumstances, to which we think a robust theory will apply: these are cases and circumstances that, while not explicitly targeted by the author, nonetheless raise similar philosophical questions as those targeted—that is, cases in which it would be reasonable to expect the scholar's logic to apply. Imagine a theorist discussing subject M decides to target cases X and Y. And imagine there is another case, Z, that is philosophically similar to X and Y. If we rightly expect the theory to apply to Z in virtue of its similarity to X and Y, then cases like Z provide an appropriate test of the theory's robustness. As we explain below, Thomas McCarthy developed a theory of reparations focused on the case of African-Americans. But we might reasonably expect such a theory to apply to other nontargeted cases that while not identical, nonetheless raise similar philosophical issues—for example, Native Americans.

We think it is useful to distinguish two different ways unrobust theories can fail to be credible, two distinct but related reasons to regard robustness as a virtue of political philosophy. The first is that a theory can be credible under some circumstances but not in others to which it should apply (these cases are sufficiently

³ Depending on their project, philosophers of science have described other relevant senses or faces of robustness (Calcott 2011; Raerinne 2013).

distinct to generate this discordant outcome). The second is that a theory that appears credible can, upon further analysis of the relevant cases and assumptions, be shown to be wholly unconvincing, to lack, what we call, *baseline credibility*. We discuss these two manifestations of a lack of robustness in turn.

Political philosophers can discuss very narrow subjects. But they typically develop reasonably general theories about broad subjects—from the nature of sexism and racism, to the conduct of a just war, the wrong of domination, the rules governing international migration and the organization of ideal societies. For completely understandable reasons, theorists are rarely able to discuss every case and circumstance to which their theory applies. Accordingly, theorists generally discuss a few key cases and focus on a specific, simplified set of circumstances. But as with empirical scholarship that focuses on a limited set of cases or uses simplifying assumptions, there is a persistent risk in political philosophy that one's theory will be problematically ad hoc; built by capturing how a subject does or should work under a limited set of conditions, the theory's success may be contingent on those conditions, failing to remain credible when tested against the broader range of conditions and cases the author is targeting.

Imagine a theory about some important subject, Theory 1, that targets conditions A, B, and C. Now imagine that Theory 1 includes an explicit and credible account of the subject under condition A. If Theory 1 does not hold in some or all of the other conditions which the author targets but may not have discussed—that is, conditions B and C—then the theory lacks robustness—it warrants credence in some cases (A), but not in others (B and C). Presumably, there is some crucial difference that sustains the theory's credibility with respect to A, but not to B and C. But assuming that the theory captures what is occurring in conditions B and C will necessarily lead one off course, causing the theorist and her readers to credit theories that fail in cases of interest and matters of substantial import. Ronald Dworkin, for instance, argued that harming others via market transactions, actions resulting in the harmful raising or lowering of prices, was permissible. He used the example of a swimming competition (Condition A), in which the competitors self-consciously and permissibly set back their rivals, to argue that harms imposed in markets (Condition B) were permissible (Dworkin 2011, 287). But Hayden Wilkinson has recently argued that while Dworkin's theory may hold for sporting competitions, it fails in a variety of real-world market situations—for example, wheat markets (Wilkinson 2022). In the global recession of 2008, for example, investors poured money into wheat, which was seen as a safe investment, sending the price of wheat skyrocketing, causing hunger and riots in many places. When wealthy people purchased wheat during that time, they caused harm. This harm, Wilkinson argues, is not akin to the harm caused by winning a swimming competition. If correct, this critique shows that Dworkin's swimming-based account of market harms lacked robustness.⁴ By contrast, robust theories remain plausible once we move

beyond our simplifying assumptions and cases—that is, the theory continues to be credible when applied to targeted conditions B and C. That is why they are good candidates for extension and development.

An implication of our argument in a case of this sort is that one might achieve robustness by narrowing one's argument, specifying that one is only targeting condition A. A can be hived off from the other cases because of philosophically relevant differences between it and the other targeted cases. And explicitly investigating why a theory is credible with respect to condition A but not B and C can lead to a fuller understanding of a theory and its contours—and, in this respect, it can improve the theory.⁵ This strategy makes sense in cases in which one is not aiming for a more general claim, such as an effort to show that harms imposed during a sporting competition were permissible. But much political philosophy, like Dworkin's argument or the cases we discuss below, targets big subjects under a variety of conditions, and restricting the scope of that kind of theory would fundamentally undermine the authors' project.

A classic and clear example of a philosophical theory that appears unrobust in this way is G.A. Cohen's short book *Why Not Socialism?* (Cohen 2009). Making the case for the application of socialist principles to modern societies, Cohen explores the organization of a camping trip. Several commentators questioned whether principles for a camping trip were actually suitable for the economic organization of large societies. Those commentators were willing to accept Cohen's principles in the conditions of a camping trip. But they concluded that Cohen's camping-based theory would not reliably apply to the alternative conditions that his theory explicitly targeted (Heath 2022; Ronzoni 2012). Assuming these concerns were valid, had one built on Cohen's theory, one's effort would be wasted, failing to shed light on the principles that should govern modern social systems. Of course, Cohen might have assured the robustness of his theory by limiting his focus to camping trips. Such a maneuver is available because camping trips are so distant from large economies. But this move would have undercut the broader rationale for engaging in his project in the first place.

The second reason robustness is a virtue of political philosophy is that the baseline credibility of robust theories does not turn to too great a degree on the choices a philosopher makes outlining and defending the theory, in particular on the subset of cases and

⁴ Clarissa Rile Hayward points out a similar limitation of Douglas McAdam's "good versus evil" model of effective disruptive action (Hayward 2020; McAdam 1996). Hayward shows that an analysis of disruptive action that is not only focused on the exemplary examples of the civil rights movement of the 1960s might see that disruptive action can also work by "interrupting privileged people's motivated ignorance." Thanks to a reviewer for pointing us to this example. See the related discussion of Rawls' account of the right to civil disobedience below.

⁵ We are grateful to one of the reviewers for pointing out that investigating why a theory cannot credibly account for certain cases is a productive way to strengthen one's theory.

assumptions used to develop the theory. This issue is closely related to the issue just described, but it is distinct from it. Recall the issue above. It entails that Theory 1, which is credible under condition A, holds under other conditions targeted by the author (e.g., B and C). Here the question is whether Theory 1 is actually credible or compelling, whether it warrants belief, even under condition A. Why might this be a problem? Because an analysis of the subject in A may not have fully and correctly identified what mattered theoretically. Theory 1's ostensible persuasiveness under A may have depended on other conditions not having been explored. Understanding that this is the case may only become clear when other targeted conditions (e.g., B and C) are analyzed. This is especially likely to be true if B and C are similar philosophically to A.

In 1966, at the height of America's civil rights era, John Rawls outlined an influential, fairness-based explanation of the right to be civil disobedient in a reasonably just democratic regime, a right to break the law via public, nonviolent, civil action aimed at persuading the majority. Notably, Rawls argued that those who undertake civil disobedience should accept legal punishment for their actions. As many commentators, such as Juliet Hooker, have pointed out, a consideration of the long history of Black activism and other instances of civil disobedience (conditions B and C, in this case) reveals a problematic narrowness in Rawls' analysis (Hooker 2016). For instance, questions of fairness arise if citizens are asked to accept punishment for their efforts to correct an injustice (e.g., Jim Crow), when those same citizens already suffer from that societal injustice. And studies have shown that disobedience might have legitimate ends other than persuading a recalcitrant majority—such as building a sense of effective political agency amongst members of a minority group (Pineda 2021). In discussing the fairness-based duties of the protestors, and the communicative force of disobedience, Rawls did not weigh the considerations just described. In our view, these concerns do not just threaten the credibility of Rawls' argument when considering a broader range of civil disobedience (Conditions B and C), they arguably mean that we cannot be confident that Rawls' argument has *baseline credibility*, credibility even in the context of the situations he discussed explicitly (Condition A). That is, his theory may be *ad hoc*, its ostensible persuasiveness dependent on a failure to explore other instances of the practice in question. A robust theory, by contrast, would not suffer from this default.

Beyond looking at circumstances explicitly targeted by the author, we can sometimes gain purchase on a theory's robustness by exploring *nontargeted* conditions (D, E) that raise philosophically similar issues to A. These cases are sufficiently similar to A that we would expect the author's theory to apply to them. If the theory fails in D or E then that might indicate the analysis of A was faulty. Alternatively, an analysis of D and E may show that the original analysis did not correctly pick out what mattered theoretically in

A. In these kinds of cases, the failure of a theory when considering alternative conditions (D and E) reveals the more general failure of the theory. And where concerns of this sort are not countered, it will be reasonable to harbor doubts about the basic credibility of the theory.

For example, Thomas McCarthy has explored the justification of reparations. His particular focus is the legacy of American slavery and the potential debts owed to African-Americans (McCarthy 2002; 2004). He offers a multifaceted argument that turns on several premises—for example, subjects of injustice are owed recompense, an original injustice can be temporally extended via discriminatory institutions and practices targeting those identified as being members of the same group as those who suffered the original injustice, participants in a corporate entity that was responsible for the original and ongoing injustices may have duties to provide the relevant recompense even if neither they nor their progenitors were authors of the original injustice. In light of these claims and a recapitulation of African-American history, McCarthy concludes that reparations are owed to African-Americans, all things considered.

In making his careful and contextual argument, McCarthy gives no consideration to other cases of historical injustice that raise similar, if not identical, philosophical issues, cases to which his premises seem very likely to apply. Among others, these include Native Americans, Native Hawaiians, and Asian-Americans. And he offers little justification for averting his analysis from these cases. We are not ill-disposed to McCarthy's conclusion. But exploring a subject by focusing on a single set of circumstances and sidestepping other cases raising philosophically similar questions leaves a theory vulnerable to concerns about its robustness. For instance, when applied to these other cases, the theory's conclusions might clash with our fixed intuitions. That would suggest that our pre-existing intuitions about the need to provide reparations to African-Americans, rather than McCarthy's justification, are actually supporting the credibility of his claims. Alternatively, we might find that his theory is overinclusive. It might apply to so many different groups that actually accepting his claims would generate unfeasible and morally unattractive demands on contemporary Americans. It is quite plausible that an analysis of additional cases might undermine McCarthy's case for reparations with respect to African Americans. By implication, McCarthy's narrow strategy gives readers reason to be concerned that his theory and conclusions only seem plausible because he focused on just one case. Readers have reason to doubt the baseline credibility of the theory because the theory may be *ad hoc*. McCarthy's theory *might* be correct, but readers are left without assurance.

Robust theories are not *ad hoc*—they are solid and not problematically contingent on the cases and assumptions on which they are built. And theories that have had their robustness explicitly tested are not subject to the same level of credibility concerns we have just discussed. Readers are assured that the

theory's persuasiveness is not dependent on avoiding problems and questions that would be raised by considering relevant alternative conditions. As a result, they can be more confident in the theory and its conclusions, that the theory actually captures something about the world, that it is worth responding to and developing.

One might have the sense that our account of robustness recapitulates, in different terms, well-known critiques of ideal theory. Partisans of nonideal theory have challenged some forms of ideal theorizing for failing to offer clear action guidance in nonideal conditions and for generating conclusions that seem not to hold in the real world (Brennan and Pettit 2007; Stemplowska and Swift 2012). Issues of robustness certainly come up in debates of that sort. In the view of certain critics, for instance, ideal theorists want to offer guidance about what is to be done under nonideal conditions, but, in our terms, their ideal theories lack robustness; the ideal theories either fail to offer credible guidance in those targeted nonideal conditions or they generate guidance that is normatively unattractive.

There is a clear overlap between the concerns driving our interest in robustness and the concerns of critics of ideal theory. But robustness is the more general topic, not reducible to the ideal/nonideal distinction. There is a simple way to illustrate the independence of robustness from the debate about the translation of ideal theory into nonideal conditions: showing that issues of robustness can arise for ideal theories without considering their application to nonideal settings.

John Rawls' stability-based critique of utilitarianism in a *Theory of Justice* illustrates the point (Rawls 1971).⁶ In our terms, though focused squarely on the ideal, Rawls offers a critique of utilitarianism's robustness by claiming that it lacks stability. Early in the work, Rawls outlines why he thinks his own principles of justice would be selected in the original position and rival views, like utilitarianism, would not. Later Rawls offers another reason to prefer his view: stability. Stability, on Rawls' account, is a too often overlooked virtue of a well-ordered a society. He worries that in otherwise ideal conditions—that is, a well-ordered society—when the particular strictures of the original position are relaxed and individuals begin to consider their own self-interest, support for a just scheme of cooperation might falter (Sirsch 2020). In a stable society, however, individuals would support the institutional setup not because of coercion or because they were angels, but because they embraced that society's principles of justice. “To insure stability,” Rawls argued, people “must have a sense of justice or a concern for those who would be disadvantaged by their defection, preferably both” (Rawls 1971, 497). Rawls contended that outside the original position, his vision of a just society would generate the requisite sense of justice. By contrast, he contended that an analysis of psychological

principles showed that an ideal utilitarian society would not generate the same reciprocity-based concerns and, as a result, the same kind of stability. Under more realistic but still ideal conditions, therefore, utilitarian principles would become less attractive than they appeared in the original position. In this respect, utilitarian theories lacked robustness.

Our claim is not that all ideal theories are subject to worries about robustness. For instance, theories narrowly aiming to defend certain principles or an approach to theorizing without considering any facts or cases may be immune from these concerns. But since many ideal theories, particularly those focused on institutions, like Rawls', incorporate empirical assumptions and facts in their elaboration, concerns about the robustness of a theory may arise without considering nonideal conditions (Sirsch 2020, 40). The reasons to consider robustness a virtue, the import of retaining credibility across a variety of targeted circumstances and not crediting unwarranted conclusions, may often apply to ideal theories whether or not they are applied to nonideal circumstances. Moreover, we think it requires little argument to realize that nonideal theories, like the theory of reparations we discussed above, raise issues of robustness without those theories engaging at all with the implications of ideal theory. If what we have argued is correct, then concerns about robustness are not identical to concerns about the application of ideal theories to nonideal conditions.

TESTING FOR ROBUSTNESS

Robust theories are more credible than those that lack robustness. By implication, political theorists should take the robustness of their theories seriously. Moreover, as we have intimated in the discussion above, one has reason for concern even in cases of reasonable uncertainty, in cases in which it would be sensible to question a theory's robustness (even if, by chance, the theory happened to be robust). When the costs of putting faith in theories that lack robustness are significant, as we believe them to be, then it will often be reasonable for theorists to show that their work is robust. Doing so involves explicitly testing for robustness. In this section, we offer two examples of how to test for robustness, both inspired by practices in the social sciences.

The first test is the *alternative assumptions test*. An empirical analyst might assume that the characteristic being studied is distributed normally. Having done so, she might come to particular conclusion about her data. To try to ensure she is not being led astray, she might test whether an alternative and equally plausible distribution would result in a very different conclusion. Refraining from doing so, the author leaves room for her audience to doubt that her finding was a plausible description of the world rather than an artifact of her assumptions. A similar logic applies to political theories. Political theorists make assumptions about their topic and their theories turn on those assumptions. These include assumptions about human nature,

⁶ Rawls elaborates his understanding of stability in later works, but we do not think the changes he makes alter the general force of our comments.

individuals' taste for risk, the likelihood that a particular institutional setup will engender political stability and the degree to which relations among friends or family members approximate relations among members of a political association. Frequently, these are not the only plausible assumptions the theorist could make. Robust theories still succeed when such assumptions are altered in reasonable ways. If a theory fails this test, then the theory appears *ad hoc*.

But what do we mean by altered in reasonable ways? Within what universe of assumptions must a theory be robust? These are crucial questions for our argument. Surely, not every assumption must be investigated. Some assumptions are shared by all the participants in the relevant scholarly debate and can be treated as settled. We think it is conceptually useful to restate the question: what are the alternative assumptions that, when tested, are most likely to cause us to significantly increase or decrease our estimation of a theory's credibility? The most obvious alternative assumptions meeting this standard are those that: (a) would be important to the theory, and; (b) compared with the original assumptions, just as closely approximate or more closely approximate some or all of the situations targeted by the author. If a theory is about a certain subject in situations S and T, then the theory should be robust to assumptions that likely characterize S and T as well as or better than the assumptions the author employs to make her case. An implication of this standard is that an author may preempt a demand for testing of an assumption by showing that plausible alternative assumptions are actually less likely to characterize the situations she targets than the assumptions she uses.

By testing high-leverage alternative assumptions, an author sheds light on the credibility of her theory. A theory's failure when tested in this way can suggest different things. When the alternative assumption better characterizes only part of the range of situations targeted by the author (e.g., only situations like S), a failure indicates that the theory reliably applies only to the subset of situations better covered by the author's initial assumption (i.e., situations like T). In contrast, when the alternative assumption better characterizes all or almost all of the situations targeted by the author (e.g., situations S and T), the failure indicates that the theory is largely misjudged, that it lacks baseline credibility.

An essay by Zach Barnett neatly illustrates the utility and import of the alternative assumptions test (though Barnett does not use this terminology) (Barnett 2020). Barnett critiques a well-known theory, defended by Jason Brennan among others, that it is irrational to vote in the name of changing the outcome of an election or referendum (Brennan 2012). Defenders of this view standardly assume that voters' behavior during an election is well enough captured by a binomial model of vote distribution. Using a binomial model, in cases where one candidate is favored to win, say 51% to 49%, each voter's decision is modeled as a biased coin flip with a 51% chance of selecting the favored candidate. Of course, if enough biased coins are flipped, there is an

extraordinarily high likelihood that they will be distributed consistent with their bias—51% versus 49%. By implication, given a reasonably large population, assuming a binomial model for voting will make it seem extraordinarily unlikely that one vote will be decisive, even in what appears to be a reasonably tight election. Clearly, this assumption matters to the theory. But do we have reason to consider an alternative?

Barnett notes that assuming a binomial model of voting effectively precludes the possibility of upsets (again, if enough coins are flipped, the outcome will precisely reflect the bias of the coins). But political upsets do occur. Accordingly, Barnett considers an alternative assumption about the distribution of votes, one that more closely approximates the situations targeted by these theories. And Barnett illustrates that if we alter those modeling assumptions in realistic ways, we get a very different result. What are the novel assumptions? That upsets are possible—that is, both candidates have at least a 10% likelihood of winning—and that outcomes will cluster around what the polling suggests the outcome will be. Under these conditions, the likelihood of one's vote being decisive is much higher than the standard critique allows. And because Barnett's assumptions are more plausible than those employed by the standard critique (according to which we should almost never observe any upsets), Barnett concludes that the standard approach to the rationality of voting is deficient.

This example shows exactly how employing the alternative assumptions test can keep philosophers from engendering misunderstanding. Had theorists and philosophers questioned their assumptions about the distribution of votes, the embrace of unrobust theories might have been avoided. Just as importantly, Barnett's essay well-illustrates that philosophy is the work of many hands, we need not solely rely on authors themselves to test the assumptions of various theories but can and should embrace this kind of test as a valuable element in the practice of political philosophy.⁷

Here's the second test we propose: the *alternative cases test*. This test is inspired by political scientists', economists', and sociologists' concerns about external validity. The examination of alternative cases can help determine whether causal claims that find support in one case or a set of cases remain persuasive in other cases of interest. For example, a theory of the behavioral response to increased savings which is tested via an experiment in one village may fail an external validity test if its finding does not hold in other villages or in larger communities in which we would expect the theory to hold (Deaton 2010, 448). Similarly, a theory of American electoral behavior tested on data from local elections might have its external validity investigated by examining data on national elections. Presumably, a theory of American electoral behavior targets

⁷ We thank one of the reviewers for suggesting the idea of a scholarly division of labor.

both kinds of cases and should, therefore, succeed in both sets of cases.

For political theories, an examination of alternative cases can be used to test a theory's robustness in both senses outlined above. If the cases in which the theory fails are sufficiently distinct from those in which it succeeds, we might infer that the theory is not credible when applied to some of the cases that the author targeted (the difference in the cases likely explains why the theory works in one set of cases but not the other). Or, if the two sets of cases are sufficiently similar, we can infer that the alternative cases test reveals a fundamental problem with the theory's baseline credibility.

As with alternative assumptions, the universe of cases against which we can test theories reasonably are those that would likely cause us to increase or decrease our estimation of the theory. The most prominent set of cases are those targeted by the author. One successful example of a theorist who tries to show that her theory satisfies that alternative cases test is Anna Stilz, in her examination of states' claims to their territory (Stilz 2011). Stilz's theory does not target a small subset of states. With a broad focus, she develops a theory of what, if anything, gives a state the right to the territory within its boundaries. Within the constraints of an article-length work, Stilz shows that the theory functions across different conditions—situations in which states are ideal, are not ideal but respect human rights and are simply incapable of advancing and defending individuals' rights. She also illustrates the theory's strength across divergent cases, from the United States, to Somalia, to occupied Germany. Stilz does not use the terms robust or robustness, but by carefully showing that her theory succeeds in a variety of cases, she illustrates that the theory is not only deductively sound but can be relied on as an explanation of states' rights in a wide variety of situations.

As we noted earlier, cases targeted by the author are not the only cases that can shake or build our faith in a theory's robustness. Empirical theories about certain causal mechanisms can be undermined when they are tested against cases that the author did not target. Generally speaking, those are cases in which we expect to observe the same or similar causal dynamics as in the cases actually targeted by the author. With respect to a political theory, cases that raise similar philosophical problems to those targeted by an author can also provide a good test of that theory. A failure of the theory in such cases indicates that the theory is mistaken. We discussed a scenario like this earlier in this essay—that is, McCarthy's theory of reparations for African Americans. Similarly, imagine that a philosopher worked out a general theory of justified self-defense and mainly tested that theory in the context of interstate wars. An alternative cases test of such a theory might examine its application to philosophically similar conflicts that are not interstate wars. If the theory of self-defense failed in cases the philosopher had not targeted, but which were philosophically similar nonetheless, we would have good reason to doubt the theory's credibility.

We have outlined two different kinds of tests a theorist (or her critics) might undertake to ensure that her theory is robust. But how much testing must she do? Assessing this requires a measure of judgment. A theorist cannot assess every single assumption or case in the name of increasing the credibility of her theory. And there are undoubtedly diminishing returns to this kind of testing. Stilz does not discuss all the cases targeted by her theory. She discusses a variety of distinct scenarios and focuses on those that seem most likely to upend her work. We think that strategy makes good sense. Consistent with our focus on what can cause us to doubt the credibility of a theory, authors should prioritize cases that are most likely to dent their theory's credibility. The author and the reader may have little cause for concern if a theory fails in a single, unimportant case. But if a theory fails in a wide variety of cases or in cases of special import, then that theory is insufficiently robust to be credible. Political philosophers should guard against that possibility.

One might ask: how distinct are these alternative assumptions and cases tests? In our view, they are very closely related, different faces of the same coin. One might for example, redescribe a test of an assumption (e.g., does a theory of state territory remain robust when not assuming that every state respects human rights?), by asking whether a theory remains credible when applied to cases characterized by that assumption (e.g., does a theory of state territory remain robust when discussing specific states that do not respect human rights?). Still, we think it is useful to distinguish between these tests. For instance, testing alternative cases might be called for when one is not quite sure which are the most important assumptions to test; the examination of alternative cases may reveal which assumptions ought to be questioned. And, as we explain in the section below, there might be instances in which a test of alternative assumptions reveals more significant inadequacies than an examination of a few, potentially exceptional, cases. In sum, both tools are useful and both are aimed at understanding the same thing: whether theories can withstand reasonable changes to their architecture; whether they are problematically ad hoc.

Notwithstanding our arguments, there may be instances in which one simply cannot or should not undertake the tests we have outlined—situations in which authors face constraints imposed by limits of time, of word count, or specialization come to mind. There can be, we acknowledge, a kind of division of labor in political philosophy—one theorist might outline a theory in one article, focusing on ensuring its internal consistency, an essential virtue of political theory, without testing its robustness, knowing that if the theory appears credible initially its robustness can be examined in other, later work, perhaps by other scholars. Still, when not testing for robustness might be reasonable, one should temper one's claims accordingly. One should not assume one's theory is robust. And one might acknowledge this explicitly, while pointing out areas for further inquiry. These efforts help ensure that authors and readers do not mistakenly

assume that theories apply more broadly than is justified or give credence to theories that are fundamentally flawed. Political theorists and philosophers of every possible stripe should enthusiastically seek to bypass this avoidable misunderstanding.

MULTICULTURALISM AND THE ALTERNATIVE ASSUMPTIONS TEST

We claim that our tests help clarify whether theoretical arguments are sufficiently robust. To show this, we apply them to a well-known theory of multiculturalism in this section, and an important theory about migration in the next section. While there are many arguments for cultural rights, we focus on Will Kymlicka's, the best-known version, illustrating why it appears unrobust.

In a pair of prominent books, Kymlicka argued that minority cultures have certain rights and those rights warranted state support for those cultures (Kymlicka 1989; 1995). Cultural membership was such an important good that liberal political theory was amiss in overlooking it. Liberal theory, therefore, required amendment.

Kymlicka's argument had several bases. Kymlicka tied cultural membership to liberal autonomy, arguing that people must be situated within a culture to make meaningful autonomous choices. Freedom and choice, on Kymlicka's view, were not particularly significant in a vacuum. "Whether or not a course of action has any significance for us depends on whether, and how, our language renders vivid to us the point of that activity. And the way in which language renders vivid these activities is shaped by our history, our traditions and conventions" (Kymlicka 1995, 83). It is through our cultural narratives that we can make intelligent judgments about how to lead our lives. "In this sense, our culture not only provides options, it also provides the spectacles through which we identify experiences as valuable" (Kymlicka 1995, 83). A culture allows individuals to make sense of who they are, the possible roles they can play and the choices they can pursue. Giving meaning to our options and to the practices around us, culture gives substance to freedom. But while vibrant majority cultures require little support to contribute to valuable choice-making, Kymlicka contended, minority cultures might fade without support from the state.

On its own, an argument about the import of culture to choice would not suffice to sustain an argument for state support. One might simply note that members of one pressured culture could choose an alternative. Accordingly, Kymlicka bolstered his theory, claiming that culture switching in the face of cultural collapse undermined self-respect: "people's self-respect is bound up with the esteem in which their national group is held. If a culture is not generally respected, then the dignity and self-respect of its members will also be threatened" (Kymlicka 1995, 89; 1989, 165). Here, Kymlicka recalled the argument of John Rawls, who held that self-respect was a primary good—a good that all citizens need if they want to pursue their plans and

projects. If the ability to exercise autonomy depends on joining a dominant culture, members of nondominant cultures cannot regard themselves as full members of that community, they lack the respect of others, their self-respect is, in turn, undermined. Sapped of respect, they will be less likely to pursue their plans and projects, they will be cut off from their vision of the good life. Switching cultures is hard: leaving one's culture, Kymlicka says, "is best seen as renouncing something to which one is reasonably entitled" (Kymlicka 1995, 86).

Finally, Kymlicka added another argument, one based on fairness and equality (Kymlicka 1995, 108–15). The dominant national community often gets for "free" what minority communities must struggle to maintain. For example, a state must conduct its business in a language (or two or three), which privileges some language communities but not others. Since language and culture are tied together so closely, this means that many states inevitably favor some cultural groups over others. Members of favored groups have richer and more frequent opportunities to exercise their agency and to profit from self-respect. Disfavored linguistic groups are bound to have a hard time surviving. This state of affairs seems to violate notions of equality and fairness. State support for minority cultures, Kymlicka argued, is then often needed to repair this imbalance, ensuring that members of those cultures have options consistent with the respect they are owed—everyone, Kymlicka argues, should have a fair opportunity to be free within their own culture.

But which cultures are owed support? There are many plausible groups that might be good candidates for state funding, perhaps dozens or even more in some states. Drawing on Ernest Gellner's influential argument linking modernization with the nation-state (Gellner 1983), Kymlicka holds that "modernization involves the diffusion throughout a society of a common culture, including a standardized language, embodied in a common economic, political and educational institutions" (Kymlicka 1995, 76). On this view, national cultures have social, political, and economic institutions, which allow them to give people a variety of options from which to choose. National cultures—which Kymlicka uses interchangeably with societal cultures—are large enough so that one can lead a fully modern life within them. One can go to school in the national culture's language, attend a university and get one of many possible jobs within the national culture.

Of course, many cultural groups are not national cultures. Not all immigrants, for example, bring a national culture with them when they move from one country to another, though they might carry specific cultural practices. Kymlicka does not ignore immigrants and other groups. Instead, he thinks "polyethnic groups"—immigrants, refugees, and some religious groups—can integrate into larger national cultures without threat to their self-respect (Kymlicka 1995, 30–1). Unlike national cultures, members of other groups must generally follow the laws of the larger society. They should not claim special protections for their group, though they can expect that their religious

and ethnic traditions are to be respected, which may mean an exemption from certain laws for individuals, as long as the relevant practices do not violate liberal rights. It may mean, for example, that a Canadian Mounty who is a Sikh need not wear the traditional Mounty hat, but such an exemption does not grant rights to Sikhs as a group, nor does this exemption undermine anyone's rights. Further, members of these groups should not expect to receive extra support for their language or anything else that might inhibit integration into their new national culture.

We think Kymlicka's argument could have been made more robust, and more credible as a result, had he employed the alternative assumptions test. In particular, he ought to have considered what it would mean for his argument if national cultures were not the *only* cultures with the capacity to frame choices and generate the valuable forms of self-respect on which he bases his argument. Kymlicka assumes that the unique character of national cultures is of critical import. But if it is the case that nonnational, nonsocietal cultures can give shape to individuals' choices and undergird their self-respect, then national cultures are not the only game in town and the credibility of Kymlicka's case for supporting minority national cultures is substantially weakened. Why spend collective resources defending weakening national cultures when those resources might be employed to achieve more pressing ends?

This assumption matters: Kymlicka's theory seems to depend on it. But is there a plausible and clear alternative, one that just as likely characterizes the situations Kymlicka worked through? We think so. What if individuals can be members of multiple overlapping cultures, cultures which frame their choices. And what if they could derive self-respect from a range of them? African-Americans, for example, are neither a poly-ethnic immigrant group nor a national cultural group. Yet, surely, African-Americans' choices are impacted by that culture, and they derive self-respect from it, as well as from America's national culture, their religious communities and so forth. While Kymlicka calls the situation of African Americans "distinct" these complications are not confined to the case of African Americans (Kymlicka 1995, 24). The choice-making context and self-respect of a person of Indian descent in Britain is undoubtedly a contingent matter. They might be formed by her connection to India, her religious group, the UK, her ethnic group in India, the British Asian community, or some combination of these groups. We think it is likely impossible for anyone to posit which group is most important for this person's self-respect, nor can we say that one culture dominates her context of choice. The same might well be true of a farmer from Nebraska or a musician in Miami or an Ultra-Orthodox Jewish person in New York, though they ostensibly share a national culture. Here's the upshot. It is plausible that national cultures are not the only source of individuals' self-respect. If true, that reality would undermine Kymlicka's argument for state support of national cultures. By implication, Kymlicka's theory plausibly lacks robustness. Moreover, a credible liberal theory of culture based in self-respect

would likely need to recognize, as some of our examples illustrate, the important possibility that for those who face discrimination and oppression, a national culture may *undermine* self-respect, while other cultural contexts sustain it.

Our argument is not that national identity cannot have an important place in liberal theory. Our claim is that Kymlicka's theory stands on a questionable assumption. Once we put Kymlicka's argument to the alternative assumptions test, the argument seems to lack credibility. More work and words would be needed to show this definitively, but there seem to be too many kinds of groups, and too many possible sources of self-respect for Kymlicka's argument to sidestep the concerns we have raised. Had Kymlicka explicitly weighed our alternative assumption perhaps he might have been able to show why self-respect still warranted support for minority national cultures. Alternatively, he might have concluded, as have others, that if national identity is to have a place within liberalism it may need a different foundation than self-respect (Carens 2000; Patten 2014). Regardless, our analysis illustrates the value of assessing one's own work using a test of this kind.

In the section above, we claimed that the alternative assumptions test and the alternative cases test overlap. Still, we think Kymlicka's argument helpfully shows why it is useful to distinguish between these two ways of testing a theory's robustness, useful because it can cause both authors and readers to reflect on different dimensions of robustness.

Kymlicka, for instance, discusses a variety of alternative cases, but they do not lead him to reassess his assumption about national cultures as the exclusive source of cultural self-respect. Kymlicka's main examples of immigrants are Americans immigrating to Sweden. Americans in Sweden, Kymlicka argues, went voluntarily, and volunteer immigrants relinquish their rights to their national culture, pushing the issue of self-respect aside for immigrants (Kymlicka 1995, 96). Yet this raises the issue of self-respect for other kinds of groups—people who are not obviously voluntary immigrants (like those who move because of poverty or dangerous living condition), or African Americans, or very small societal cultures that are hard to sustain, among others. To his credit, Kymlicka mentions several of these cases, groups whose cultural claims cannot be so easily dismissed as those of Americans in Sweden. As he notes there are "many such hard cases and gray areas" and he concluded that his theory could not be expected to resolve these exceptions (Kymlicka 1995, 101). Uncharitably, one might say that if one finds many targeted cases to be exceptional, then, perhaps, one's theory lacks robustness. But, more importantly, labeling many cases as exceptions is not a move available if one is considering alternative assumptions. And though we cannot be certain, had Kymlicka explicitly considered an alternative assumption about the sources of self-respect, he may have come to a different conclusion about his argument and these cases. That may well have generated an argument readers could regard as robust.

To be clear, our goal, in this limited space, has not been to offer an unanswerable critique of Kymlicka's theory. Instead, we aim to make clear the value of the alternative assumptions test and to show why Kymlicka's case would be strengthened if he anticipated objections of this sort, taking seriously the important role that these assumptions play in his theory.

CLOSED BORDERS, CLIMATE CHANGE, AND THE ALTERNATIVE CASES TEST

We now examine the robustness of a prominent investigation of the morality of borders and immigration: Michael Blake's *Justice, Migration, and Mercy* (Blake 2019). In doing so, we highlight the utility of the alternative cases test. Blake acknowledges but largely sidesteps the implications of climate change, leaving his readers with reasonable concerns about the theory's robustness. In a world where human-made climate change impacts most polities and will fundamentally impact the shape of human migration, a world like our own, even a reader who accepted Blake's theory *in toto* could not leave it sure that she possessed a reliable understanding of whether states have a right to exclude most foreigners. This finding bolsters our claim that robustness is a virtue. And it underlines the value of using the alternative cases test. If a powerful theory like Blake's could be strengthened by assessing its robustness and by confronting, rather than sidestepping, a key characteristic of the cases his theory targets, then this is very likely true of other theoretical arguments.⁸

In *Justice, Migration, and Mercy*, Blake's argues that states have a right to exclude others. For Blake, states are jurisdictional projects—states establish coercive, legal jurisdictions that put citizens into special relationships with one another and to their state (Blake 2019, 68). Individuals subject to a state's jurisdiction have a duty to establish and preserve institutions that satisfy each other's human rights. But the "fact of jurisdiction" not only generates duties, it also restricts them. Citizens have a defeasible right to avoid further obligations (Blake 2019, 78). In particular, citizens have a right to exclude outsiders from the state's jurisdiction, to keep them from entering the state and thereby generating additional duties for existing inhabitants.

Blake's focus on jurisdiction is grounded in concerns for human rights and autonomy. Our duties to those with whom we share a jurisdiction turn on those

concerns. For similar reasons, inhabitants of well-off polities have duties to proactively extend protections for human rights to those without the option of living in a rights-respecting state, such as people suffering from state persecution. In those cases, the normal justification for excluding others (not having to bear unjustified duties) will not apply since one is already bound to support those individuals' human rights. And, in Blake's view, one prominent way the well-off can satisfy these duties (though not the only way) is by allowing the immigration of those disadvantaged individuals. Beyond this select group, however, the duties of the well-off are relatively limited.

One might take issue with the logic of Blake's arguments, his conclusions and any unforeseen implications of his theory. Our focus is whether the theory is robust and, in particular, whether we can be confident it holds in cases his theory targets but largely avoids discussing, cases touched by climate change. Climate change is likely to influence international migration directly (by reducing the attractiveness of living in certain places) and indirectly (by influencing other drivers of migration like wars, poverty, and urbanization, and so on). These factors complicate the image Blake draws. Most obviously, citizens of polities where migrants would like to move (such as the United States, the United Kingdom, France, etc.) are contributing disproportionately (both as individuals and via their collective institutions) to environmental changes that may have already negatively impacted the well-being of potential migrants. Does this vitiate or outweigh the right to exclude these potential migrants? It might.

For instance, we think we have a right to exclude our next-door neighbors from living in our yards. But imagine that our powerful air conditioning systems contribute to uncomfortable and permanent increases in the temperature of their homes. Imagine further, that we have known about this process for some time. If our neighbors asked to tent in our spacious backyards, we would be wrong for refusing these requests (of course, we could pay them off, but imagine they had nowhere else to go). We have wronged our neighbors, made their homes uncomfortable places to live. We did so knowingly and without their consent. As a result, it seems plausible that our claim to keep them from our yard no longer carries the same, if any, weight. This kind of logic may well apply in the context of climate change.

Of course, our example is grossly simplistic. The moral situation generated by climate change is *so* much more complicated (Caney 2012; 2014; Draper 2022). Any individual's contribution to climate change may be negligible. A single state's contribution to climate change may be large but difficult to assess. Potential migrants may have contributed to climate change themselves. Climate change will influence the migration of individuals in a variety of different ways and even so, many potential migrants would seek to move without climate change (Lustgarten 2020). Many of those who contributed to climate change are dead. The duties to respond to an impending migratory crisis may not

⁸ *Justice, Migration, and Mercy* is an especially apt case to consider because the author has a demonstrated interest in the issues raised by our analysis. In an essay entitled, "Global Distributive Justice: Why Political Philosophy Needs Political Science," Blake contends that debates about the scope of distributive justice run aground on contestable assumptions about political institutions (Blake 2012). Participants in the debates, he argues, are not fundamentally divided by differences in deductive approach but by conflicting and unsupported assumptions with respect to the coercive character of international institutions. In our terms, Blake is noting that we simply do not know whether contributions to the debate on the global distributive justice are sufficiently robust to be persuasive.

attach most strongly to those who *caused* the change but to those who can do something about it and so on.

Here's the upshot. It is difficult to understand the implications of climate change on the right to exclude without carrying out this complicated analysis, partly theoretical and partly empirical. Blake's theory clearly targets our world, a world marked by climate change. But we do not know if the theory can withstand the challenge of climate change, or whether it needs to be changed in some way. The theory might well fail an alternative cases test, making it reasonable to be concerned that it cannot be relied on. And we think these doubts may well have been avoided were robustness a recognized virtue of philosophical theories.

To be fair, Blake does not *completely* ignore climate change. Some individuals affected by it are in the same moral position as those whose human rights are undermined by a state. "It might be the case that forms of evil other than persecution similarly demand coordinated international effort," he suggests. "Climate change, for instance, seems potentially as capable of destroying a political community as atrocity; a society that is literally underwater is, after all, unlikely to have a flourishing future" (Blake 2019, 114). This comment assimilates climate change to atrocity, the state-led violation of human rights—it is another form of evil that makes it impossible to live an autonomous, valuable life. Thinking of climate change in this way limits the implications of Blake's theory. Inhabitants of flooded countries (and others in a similar situation) must be accepted by other polities or have their human rights reliably advanced by some other means. Those whose countries are unflooded do not have the same claims. Yet, as Blake no doubt recognizes, the impacts of climate change cannot be reduced to societies that are literally underwater. By implication, extreme examples of this sort cannot be used to avoid the concerns we have raised.

Alternatively, Blake might claim that *Justice, Migration, and Mercy* presents a theory only intended to address the core case of migration justice. It does not really target cases impacted by climate change. Climate change, while obviously of great import, is an exceptional issue or one that will impact a relatively small set of cases. We do not know, however, how anyone could have any confidence that climate change will not be implicated, directly or indirectly, in what are rightly thought of as the core cases of migration or, indeed, how we could have confidence that it is not *already* implicated in the core cases of migration. Consider the dramatic flow of Syrians outside that country's borders during its civil war. It is common to assume that these movements were generated by the civil war—but some argue that droughts spurred migration directly and intensified the conflict, which, in turn, drove further migration (Kelley et al. 2015). If the effects of climate change are broadly felt, as we believe them to be, then claiming his theory does not target such cases would mean that Blake would have drastically limited the ambit of his theory. A move of this sort, of course, is not impossible. But we think it

would be inconsistent with the ambitions of this admirably ambitious work.

Justice, Migration, and Mercy features a discussion of whether Blake, a Canadian at the time, had the right to move to the United States to pursue graduate study (Blake 2019, 145). In cases of this sort, in which reasonably well-off citizens of one wealthy, geographically well-positioned polity seek to immigrate to another similar polity, climate change may not rear its head. The issues raised by that case may rightly be accounted for solely, or largely, by Blake's coercion-based theory. In cases of this sort, Blake's theory seems credible. But given the pervasive impact and the complexity of the issues raised by climate change, it is not clear how robust Blake's conclusions are. It is reasonable to fear that his theory lacks credibility.

On the second to last page of *Justice, Migration, and Mercy*, Blake notes: "Climate change, which I have not discussed in the present context, is likely to increase both economic inequality and the number of people seeking to migrate to countries of wealth" (Blake 2019, 224). If Blake is right about this, then it is difficult to avoid the conclusion that a robust account of this complex institutional issue requires an explicit treatment of the forces shaping the cases he is targeting. Were Blake to have tackled the question of robustness head-on, it might have resulted in a more nuanced or limited theory. And that theory might have been more likely to pass an alternative cases test. One would quite reasonably hesitate to board a plane in the summer if one knew its engineers only considered winter conditions. For the same reasons, one might hesitate before putting full faith in *Justice, Migration, and Mercy*.

CONCLUSION

Robustness is a virtue of the social sciences. This article illustrates why it is also a virtue of political philosophy. Like their counterparts in the social sciences, robust theories can be counted on; they work in the cases and circumstances the author is talking about. Theories that are not robust, by contrast, engender misunderstanding and cause scholars to lose time responding to and building on theories that are ad hoc. There is a solution to this problem. Theorists and philosophers can test whether their work is robust. Borrowing from political science and economics, we have outlined two straightforward ways of doing so—the alternative assumptions test and the alternative cases test. There are surely other pathways for analyzing whether and why a theory works across the cases and situations one has targeted. Still, these tests and our examples illustrate the power of this kind of analysis.

Our goal in developing a theory of robustness is not essentially critical. We think this essay outlines a constructive path for fashioning better theory. Take McCarthy's account of reparations, an account limited to a single example. Given the import of making a credible theory about an issue like reparations, we think pursuing robustness would be wise. Following our

account, McCarthy could consider how his theory would be impacted by considering alternative cases, like those of Native Americans and Chinese Americans. Such cases might pose difficult challenges for his account. But they might, for that very same reason, generate confidence that his conclusions held some real force, that they were not mere artifacts of the powerful, but solitary example he chose to discuss. Similarly, McCarthy might weigh the existence of alternatives to his central assumptions—for instance, that the sole or main duty of those who are in position to contribute to reparations is their responsibility to repair *historical* injustices. Surely those same individuals also have duties of fairness to respond to contemporary injustices, whether or not those injustices were attributable to long past, but grave, crimes like slavery. And these duties might be seen to conflict. Illustrating why this is or is not the case would better position McCarthy to offer a more nuanced and more credible account of what is likely owed to those whose lives are set back by American chattel slavery.

Buildings and planes must withstand a variety of different conditions. Accordingly, engineers take robustness seriously. If they ignore it, buildings will fall and planes will crash. The aims of engineers and political philosophers are not the same, of course. But political philosophers commonly develop theories about issues of great import in a wide variety of circumstances. And they want those theories to hold up, to withstand challenge. Accordingly, philosophers also have reasons to care about the robustness of what they build. If that is correct, theorists and philosophers should acknowledge the virtue of robustness and pursue it. Such efforts will result in less confusion and more insight.

ACKNOWLEDGMENTS

We would like to thank Alex Oprea, Claudio López-Guerra, and three anonymous reviewers for their helpful comments on earlier drafts of this article. The article also benefitted from comments from workshop audiences at Duke, the University of College London, Erasmus University Rotterdam, Yale, SciencesPo, the École des Hautes Études en Sciences Sociales, and the University of California, San Diego. In particular, the authors would like to thank Luc Foisneau, Tom Theuns, and Armando Perez-Gea for their insightful responses to our article at these workshops.

CONFLICT OF INTERESTS

The authors declare no ethical issues or conflicts of interest in this research.

ETHICAL STANDARDS

The authors affirm this research did not involve human subjects.

REFERENCES

- Anderson, Elizabeth S. 1999. "What is the Point of Equality?" *Ethics* 109 (2): 287–337.
- Barnett, Zach. 2020. "Why You Should Vote to Change the Outcome." *Philosophy & Public Affairs* 48 (4): 422–46.
- Blake, Michael. 2012. "Global Distributive Justice: Why Political Philosophy Needs Political Science." *Annual Review of Political Science* 15: 121–36.
- Blake, Michael. 2019. *Justice, Migration, and Mercy*. Oxford: Oxford University Press.
- Brennan, Jason. 2012. *The Ethics of Voting*. Princeton, NJ: Princeton University Press.
- Brennan, Geoffrey, and Philip Pettit. 2007. "The Feasibility Issue." In *The Oxford Handbook of Contemporary Philosophy*, eds. Frank Jackson and Michael Smith, 258–80. Oxford: Oxford University Press.
- Calcott, Brett. 2011. "Wimsatt and the Robustness Family: Review of Wimsatt: Re-Engineering Philosophy for Limited Beings." *Biology & Philosophy* 26 (2): 281–93.
- Caney, Simon. 2012. "Just Emissions." *Philosophy & Public Affairs* 40 (4): 255–300.
- Caney, Simon. 2014. "Climate Change, Intergenerational Equity and the Social Discount Rate." *Politics, Philosophy & Economics* 13 (4): 320–42.
- Carens, Joseph. 2000. *Culture, Citizenship and Community: A Contextual Exploration of Justice as Evenhandedness*. Oxford: Oxford University Press.
- Carens, Joseph. 2013. *The Ethics of Immigration*. Oxford: Oxford University Press.
- Cohen, G. A. 2009. *Why Not Socialism*. Princeton, NJ: Princeton University Press.
- Deaton, Angus. 2010. "Instruments, Randomization, and Learning About Development." *Journal of Economic Literature* 48 (2): 424–55.
- Douglas, Heather. 2013. "The Value of Cognitive Values." *Philosophy of Science* 80 (5): 796–806.
- Draper, Jamie. 2022. "Labor Migration and Climate Change Adaptation." *American Political Science Review* 116 (3): 1012–24.
- Dworkin, Ronald. 2011. *Justice for Hedgehogs*. Cambridge, MA: Harvard University Press.
- Editors, APSR. 2022. "Notes From the Editors." *American Political Science Review* 116 (1): v–viii.
- Fricker, Miranda. 2007. *Epistemic Injustice*. Oxford: Oxford University Press.
- Gellner, Ernest. 1983. *Nations and Nationalism*. Ithaca, NY: Cornell University Press.
- Hayward, Clarissa Rile. 2020. "Disruption: What Is It Good For?" *The Journal of Politics* 82 (2): 448–59.
- Heath, Joseph. 2022. *Cooperation and Social Justice*. Toronto: University of Toronto Press.
- Hooker, Juliet. 2016. "Black Lives Matter and the Paradoxes of U.S. Black Politics." *Political Theory* 44 (4): 448–69.
- Kelley, Colin P., Shahrzad Mohtadi, Mark A. Cane, Richard Seager, and Yochanan Kushnir. 2015. "Climate Change in the Fertile Crescent and Implications of the Recent Syrian Drought." *Proceedings of the National Academy of Sciences* 112 (11): 3241–6.
- Kuhn, Thomas. 2013. "Objectivity, Value Judgment and Theory Choice." In *Arguing About Science*, eds. Alexander Bird and James Ladyman, 74–86. New York: Routledge.
- Kymlicka, Will. 1989. *Liberalism, Community and Culture*. Oxford: Oxford University Press.
- Kymlicka, Will. 1995. *Multicultural Citizenship: A Liberal Theory of Minority Rights*. Oxford: Oxford University Press.
- Levins, Richard. 1966. "The Strategy of Model Building in Population Biology." *American Scientist* 54 (4): 421–31.
- Lustgarten, Abrahm. 2020. "The Great Climate Migration." *New York Times Magazine*, July 23. <https://www.nytimes.com/interactive/2020/07/23/magazine/climate-migration.html>.
- McAdam, Doug. 1996. "The Framing Function of Movement Tactics: Strategic Dramaturgy in the American Civil Rights Movement." In *Comparative Perspectives on Social Movements: Political Opportunities, Mobilizing Structures, and Cultural Framings*, eds. Doug McAdam, John McCarthy, and Mayer Zeld, 338–56. New York: Cambridge University Press.

- McCarthy, Thomas. 2002. "Vergangenheitsbewältigung in the USA: On the Politics of the Memory of Slavery." *Political Theory* 30 (5): 623–48.
- McCarthy, Thomas. 2004. "Coming to Terms With Our Past, Part II: On the Morality and Politics of Reparations for Slavery." *Political Theory* 32 (6): 750–72.
- Müller, Jan-Werner. 2016. *What is Populism?* Philadelphia: University of Pennsylvania Press.
- Patten, Alan. 2014. *Equal Recognition: The Moral Foundations of Minority Rights*. Princeton, NJ: Princeton University Press.
- Pineda, Erin. 2021. "Martin Luther King, Jr. and the Politics of Disobedient Civility." In *Cambridge Companion to Civil Disobedience*, ed. William Scheuerman, 56–79. Cambridge: Cambridge University Press.
- Raerinne, Jani. 2013. "Robustness and Sensitivity of Biological Models." *Philosophical Studies* 166 (2): 285–303.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Ronzoni, Miriam. 2012. "Life is Not a Camping Trip – on the Desirability of Cohenite Socialism." *Politics, Philosophy and Economics* 11 (2): 171–85.
- Shelby, Tommie. 2016. *Dark Ghettos: Injustice, Dissent, and Reform*. Cambridge, MA: Harvard University Press.
- Sirsch, Jürgen. 2020. *Designing Realistic Utopia: Ideal Theory in Practical Political Philosophy*. Baden-Baden, Germany: Nomos Verlag.
- Stemplowska, Zofia, and Adam Swift. 2012. "Ideal and Nonideal Theory." In *The Oxford Handbook of Political Philosophy*, ed. David Estlund, 373–89. Oxford: Oxford University Press.
- Stilz, Anna. 2011. "Nations, States, and Territory." *Ethics* 121 (3): 572–601.
- Urbinati, Nadia. 2014. *Democracy Disfigured*. Cambridge, MA: Harvard University Press.
- Wilkinson, Hayden. 2022. "Market Harms and Market Benefits." *Philosophy & Public Affairs* 50 (2): 202–38.
- Wimsatt, William C. 2007. *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge, MA: Harvard University Press.
- Young, Iris Marion. 2011. *Responsibility for Justice*. Oxford: Oxford University Press.