

On Seeking the Truth (Whatever That Is) through Democracy: Estlund's Case for the Qualified Epistemic Claim*

Gerald Gaus

In *Democratic Authority* David Estlund advances what I call the “Qualified Epistemic Claim,” according to which democracy is the best epistemic strategy for revealing the truth from among those strategies that are generally acceptable. I examine this claim, focusing on three issues: the scope of truth claims for which it applies, the test by which we know whether democracy is best, and the mechanism by which democracy has this supposed benefit. I conclude that although Estlund has presented some suggestive considerations in favor of the Qualified Epistemic Claim, he does not present a convincing case for it.

I. DEMOCRACY AND TRUTH

John Stuart Mill saw the public as a “collective mediocrity.”¹ Mill, and many of those attracted to popular government, have been caught in a dilemma: while convinced that popular control of governors is necessary to a free and nonoppressive society, they thought that popular sovereignty threatens rule by a “low grade of intelligence.”² Thus, for Mill, a representative government must build in a system of jurisdictions according to which those with expertise are accorded decision-making power over different facets of public (and indeed private)

* My thanks to Fred D’Agostino, Bruce Brower, David Copp, David Estlund, Kelly Gaus, David Miller, and Henry Richardson, as well as *Ethics*’ readers, for conversations, comments, and a bountiful supply of objections.

1. John Stuart Mill, “Thoughts on Parliamentary Reform,” in *The Collected Works of John Stuart Mill*, ed. J. M. Robson (Toronto: University of Toronto Press, 1977), 19:268.

2. John Stuart Mill, *Considerations on Representative Government*, in *The Collected Works of John Stuart Mill*, 19:448.

Ethics 121 (January 2011): 270–300

© 2011 by The University of Chicago. All rights reserved. 0014-1704/2011/12102-0002\$10.00

life.³ The result is a complex structure, balancing and hedging popular control with the authority of expertise.

Recent democratic theory, largely spurred by the path-breaking work of David Estlund, has replied to doubts about the wisdom of popular government by developing “epistemic” theories of democracy—seeking to show that democratic institutions can be justified at least partly by appeal to their tendency to track the truth. It is hard to overestimate how important the success of this project would be for democratic theory. No longer would democracy be an uneasy mix of fairness and stupidity. Democrats would be in a much stronger position to reply to contemporary skeptics who loudly proclaim that the people are a herd of sheep who consistently produce stupid policies.⁴ Over the past fifteen years Estlund has developed what he calls “epistemic proceduralism,” according to which democracy is a better-than-random way of arriving at the correct view of what justice requires.⁵ Democracy, he argues, is, among those procedures acceptable to all “qualified” points of view, epistemically the best. Estlund claims that this is the “just right” (102) weight to put on the epistemic dimension of democracy. Purely procedural (i.e., fairness) views make no claims that democratic procedures tend to track the truth, while “correctness” theories, which hold that democratic legitimacy is simply a function of democracy’s tendency to yield correct decisions, put too much weight on the epistemic dimension. So we still have a balance between a certain sort of fairness and epistemic considerations, but it is not the uncomfortable balance between fairness and stupidity: within Estlund’s balance, democracy scores well on both the fairness and the epistemic dimensions.

In his book *Democratic Authority*, Estlund endorses what I shall call the “Qualified Epistemic Claim”:

Democracy . . . [is] [1] the best epistemic strategy [2] from among those that are defensible in terms that are generally ac-

3. Piers Norris Turner has recently developed a systematic and insightful interpretation of Mill along these lines. See his “Mill’s Epistemic Liberalism” (PhD diss., University of North Carolina at Chapel Hill, 2009).

4. Bryan Caplan, *The Myth of the Rational Voter: Why Democracies Choose Bad Policies* (Princeton, NJ: Princeton University Press, 2007). The dust jacket of Caplan’s book depicts masses of sheep standing up—apparently to be counted. I seek to reply to Caplan’s onslaught in “Is the Public Incompetent? Compared to Whom? About What?” *Critical Review: A Journal of Politics and Society* 20 (2009): 291–311.

5. For an early statement, see David Estlund, “Beyond Fairness and Deliberation,” in *Deliberative Democracy*, ed. James Bohman and William Rehg (Cambridge, MA: MIT Press, 1997), 173–204. This idea is developed in his *Democratic Authority: A Philosophical Framework* (Princeton, NJ: Princeton University Press, 2008). All page references in my text refer to the latter work.

ceptable. If there are epistemically better methods, they are too controversial . . . to ground legitimately imposed law. (42; see also 77)

The Qualified Epistemic Claim refers to the “best epistemic strategy” within the acceptable set. Estlund is clear—at least in the first part of the book⁶—that the best epistemic strategy is to be understood in terms of the best way to track moral truth. So, at least as I interpret him, we can understand the Qualified Epistemic Claim as maintaining that, within some set of acceptable political procedures, democracy best tracks policies that are required or recommended by what is the moral truth on these matters.

An evaluation of Estlund’s Qualified Epistemic Claim—or indeed, almost any epistemic justification of democracy—must get clear about three matters. (1) First, we need to be clear about the scope of the truth claims which democracy is said to track. Presumably democracy is not the best, among all acceptable methods, for determining scientific claims. Does it best track the answers to all questions of moral truth? Personal virtue? Metaethics? Call this the *Question of Scope*. (2) Second, we must try to get as clear as possible about the test according to which we determine that democracy tends to track the truth. Call this the *Question of the Test*. In a complex theory such as Estlund’s, the Question of the Test has two parts. First, there is the phase of determining the set of eligible procedures: the set of procedures which pass an initial *acceptability test*. Once we identify the nature of, and rationale for, the acceptability test, we then must see what *epistemic test* is applied to the acceptable procedures such that democracy wins out (it is the “best epistemic strategy” of those in the eligible set). This, of course, is a critical investigation, for any doubts we have about the test will result in doubts about the Qualified Epistemic Claim. (3) Finally, we want to know why democracy passes the test. The Question of the Test seeks to understand the nature of the test; to be confident about the epistemic virtues of democracy we need to see the mechanism or special characteristics of democracy that allow it to excel on (or at least pass) the test. Call this third issue the *Question of Mechanism*.

If Estlund can supply clear and convincing answers to our three questions, we should embrace his Qualified Epistemic Claim, and we then can put Millian worries behind us. To the extent we are unclear about his answers to our questions, or are skeptical of the reasons Estlund gives to support his answers, the Millian worry will continue

6. On page 98 of *Democratic Authority*, there is a rather sudden switch to talk of “objectivity,” which seems a different issue since subjectivism is consistent with truth.

to gnaw. We may not be convinced that Mill was right, but neither can we be confident that democratic theory has put behind it the worry that it caters to a collective mediocrity, endorsing rule by a low grade of intelligence. I shall argue here that, alas, for all his ingenuity and argumentative élan, Estlund has not assuaged the worry that, in the end, although democracy is fair and avoids oppression, it is rule by collective stupidity.

II. THE QUESTION OF SCOPE

A. *Substantive Limits on Scope*

As I have said, although the Qualified Epistemic Claim refers to the “best epistemic strategy,” Estlund clearly thinks of the best epistemic strategy in a truth-tracking way. This, of course, is not necessary. One can talk about the best epistemic strategy in, say, a way that seeks to yield justified belief, where there may be a considerable gap between justified belief and truth.⁷ But for Estlund the touchstone of the notion of a “best epistemic strategy” is truth-trackingness. Moreover, he is explicit that truth-apt statements include moral standards for evaluating political decisions. It is, he says, “very difficult to deny” (31) “the truth tenet”; that is, “there are true (at least in the minimal sense) procedure-independent normative standards by which political decisions ought to be judged” (30). Given the truth tenet, we should, at least initially, suppose that the phrase “the best epistemic strategy” in clause 1 of the Qualified Epistemic Claim should be read as “most conducive to discovering normative truths relevant to politics.”

It would be natural to expect that a case for democracy that places so much weight on the idea of discovering normative political truths would have pretty explicit commitments about the nature of truth claims concerning normative standards. For example, one might say that moral truths relevant to politics are those that concern the protection of individual rights and interests and that there is a special reason to think that democracy, with its popular responsiveness, is a good way to track the truth about those matters. As Dewey noted, each person knows best where the shoe pinches him; so if the political right requires satisfying individual interests, we will want to know how each person thinks his interests are faring.⁸ Now one of the truly striking features of Estlund's proposal is that, in contrast to these accounts which defend the epistemic virtues of democracy by narrowly

7. Internalist accounts of justification may allow a considerable gap here. See my *Justificatory Liberalism: An Essay on Epistemology and Political Theory* (New York: Oxford University Press, 1996), chaps. 2 and 3.

8. John Dewey, *Intelligence in the Modern World: John Dewey's Philosophy*, ed. Joseph Ratner (New York: Modern Library, 1939), 402.

restricting the substantive scope of the truth claims (to, say, the protection of liberal rights), Estlund seeks to make his case for democracy robust in relation to the content of true views about politics by not restricting the scope of possible moral truths. That is, at least early on in *Democratic Authority*, he holds that his account of the epistemic virtues of democracy holds over a very wide range of moral truth, “whatever it might be” (5).

However, Estlund is not entirely consistent about how wide the intended scope of the theory is. As the book proceeds we encounter three different characterizations of the scope. Early on it is very wide indeed: it seems to include just about any normative standard that has been advanced as relevant to political decisions: justice, the good, virtue and vice (5, 24). In this maximally robust statement of the claim, a utilitarian conception of the moral truth would seem entirely admissible. This maximally wide reading is endorsed by Estlund’s claim that to deny his conception of moral truth is to embrace nihilism (25ff.). However, as the book proceeds, earlier references to “good or just” (26) are replaced by a sole concern with justice such that, when we get to considering the relevant epistemic test, the only concern is its tendency to yield just results (e.g., 114–15; see Sec. IV.D). Finally, late in the book Estlund introduces a core list of “primary bads”—“war, famine, economic collapse, political collapse, epidemic, and genocide”—that any acceptable theory of justice must see as things to be avoided (163; more on these in Sec. IV.E). But now we really do get a very significant substantive restriction. As I will argue below, it is not the case that all plausible theories of justice see avoiding these states of affairs as a primary desideratum of political right. It might greatly matter how these states of affairs were brought about, and thus evaluating procedures simply on their ability to avoid them would be rejected.

B. Conceptual Limits on Scope

Our first problem, then, is that there is considerable ambiguity about the substantive scope of the claim that democracy tends to track the truth. What about the scope of the idea of “truth?” Truth is itself a metaphysically contentious concept. To say that democracy has significant truth-revealing features “whatever” truth might be may be taken substantively or conceptually. In the latter sense, whatever your analysis of the concept of truth, democracy has the property of tending to it.

At one point Estlund defines truth in terms of a certain sort of thought independence—things are true, right, vicious, and so on regardless of “what anyone thinks about them” (5). This would seem to preclude some sorts of constructivism, but in any event his official view is rather different, defending a “minimal” conception of the

truth. "To help keep disputes about noncognitivism off to the side," Estlund writes, "it will be useful to admit a 'minimal' kind of truth (even if noncognitivists themselves might not always like the term). Let us say that a statement that ' x is F ' is true in at least the minimal sense if and only if x is indeed F " (25).

I suppose here that, despite the somewhat idiosyncratic insertion of "indeed" in both the main statements of the minimal view (25, 32), Estlund is appealing to a disquotational analysis: a truth claim is one that is subject to disquotation. It is often thought that all plausible theories of truth allow disquotation, and in that sense it is a minimal idea of what is involved in any theory of (moral) truth. The problem is that, while the minimal conception of truth is metaphysically minimalistic, it does not mean that it is robust in relation to theories of truth, that is, that it is wide in scope in relation to theories of truth. I note two objections, based on conceptions of truth that are at odds, or at least may be at odds, with it.

1. *The Objection That the Minimal Conception Is Too Broad.*—Suppose one is a moral realist of a type who believes that a moral statement is true just in case it corresponds to some mind-independent non-natural realm. For such a moral realist, not all moral statements of the form " x is F " need state moral truths. One may correctly utter " x is F " as a conventional claim in a moral discourse, but if it does not refer in the right way to the nonnatural realm, it does not state a moral truth. So our realist identifies two sets of correct propositions: (1) " x is F " is true if and only if it corresponds to a mind-independent moral reality according to which x is indeed F , and (2) " x is F " if in some context it is correct to utter that x is F but this does not correspond to a fact in the nonnatural moral realm. For example, it may be correct to say "We have a duty to obey our leader!" but this does not refer to a nonnatural moral realm. In this case the realist evaluating whether democracy tracks the truth would be concerned only with knowing whether democracy is a reliable way to track proposition 1; Estlund's question is different: whether it is a way to track propositions 1 and 2. Estlund's minimal conception of truth allows as true a claim that the realist sees as, say, simply conventional. Moreover, as we will see (Sec. III.C), at one juncture Estlund's analysis depends on the realist endorsing the minimal conception, and so allowing as true propositions which do not correspond to a mind-independent reality.

2. *The Objection That the Minimal Conception Is Too Narrow.*—Unless we supply a more complicated semantics, a simple interpretation of the disquotational analysis may eliminate as not truth centered some types of cognitivist theories, namely, those that involve indexicals. As is well known, it is not straightforward to show that disquotation is always allowed with indexicals. Suppose Estlund says "I am David." It

does not follow I am David is true (indeed, I am Jerry, so it is false).⁹ Now suppose I am an ethical relativist, and I say "Abortion is wrong." This is best read as having an implicit indexical of the sort "I see abortion as wrong." But then this does not bear disquotation, for it is not simply the case that abortion is wrong. In general, relativistic theories can be understood as involving such indexical truth claims which may not allow disquotation: "*x* is F" may be true when uttered by Alf's group but not when uttered by Betty's. This has been called "indexical relativism" or "truth relativism."¹⁰ Typically we need to invoke some semantics that allows substitution or introduces a context into the truth claim, but these semantics themselves pose problems.

This raises a question. It is not clear whether (1) the scope of Estlund's case for the Qualified Epistemic Claim includes only non-indexical cognitive theories or (2) whether he wants to commit himself to a semantics that allows disquotation for these sorts of indexicals. I shall argue later (Sec. V) that there are reasons to think he must be led toward scope 1 rather than scope 2. Estlund's suggestions about the Question of Mechanism, I shall contend, only have some force if indexical moral claims (as in some forms of moral relativism) are generally excluded from the family of theories that accept the truth. Put differently (with one exception), the conceptual scope of the Qualified Epistemic Claim does not include forms of moral cognitivism that employ indexicals.

It may seem that worrying about the scope of concepts of truth that Estlund can accommodate in his theory is irrelevant to political theory. I do not think so: political philosophers are often rather free and easy with claims about truth, but what qualifies as a moral truth on one view does not on another. I believe, for example, that it is wrong to claim that theories of political ethics that advance claims that do not allow disquotation do not have a conception of the moral truth. One does not have to be a thoroughgoing relativist to think that true moral claims may be indexed to speakers. Kurt Baier was, in my view, entirely correct when he insisted that "there is no a priori reason to assume that there is only one true morality. There are many moralities, and of these a large number may happen to pass the test which moralities must pass in order to be called true. It would, therefore, be better to speak of 'a true morality' or of 'true moralities' than

9. See Jesse Prinz, *The Emotional Construction of Morals* (Oxford: Oxford University Press, 2008), 6–7.

10. See Max Kölbel, "Faultless Disagreement," *Proceedings of the Aristotelian Society* 104 (2004): 53–73; Prinz, *The Emotional Construction of Morals*, 180–83.

of 'true morality'."¹¹ Different groups, Baier argued, have different true moralities. P. F. Strawson agrees: "There is no reason why a system of moral demands characteristic of one community should, or even could, be found in every other."¹² As I will argue, Estlund's support for the Qualified Epistemic Claim has plausibility under only one form of relativism, thus making the analysis more restricted than it first appears.

III. THE QUESTION OF THE TEST, PART 1: IDENTIFYING THE ACCEPTABLE SET OF PROCEDURES

A. *The Contenders*

Thus far I have simply tried to determine what theories of political-moral truth are within the ambit of the Qualified Epistemic Claim. Unless we know the scope of truth claims—in both the substantive and conceptual senses—we just do not know what the Qualified Epistemic Claim says. I have tried to stress that scope is more restricted than it first appears. This is no deep criticism: some restriction on the scope is necessary, though I will argue that the restriction is more controversial than one may initially think. Having attained some traction on the scope of truth claims, let us turn to the Question of Test, the criteria according to which democracy has the advantages over competitors. As I said, in a qualified view such as Estlund's, the Question of Test has two stages: first identifying a set of eligible procedures and then showing that within this eligible set democratic procedures are "best." I consider the first stage in this section.

As I understand it, Estlund's "Qualified Acceptability Requirement" identifies some justificatory public *P* composed of "qualified points of view" (45). Only if some procedure *S* can be accepted by all the qualified points of view is *S* eligible, and so will be allowed to advance to the second, epistemic, stage of the test. Reasonably, Estlund's discussion is devoted to the merits of four different ways of characterizing how the first, eligibility, decision is to be made. They are:

Actual Acceptance as a Necessity: A procedure *S* is eligible only if *S* is accepted by all in the political community.

11. Kurt Baier, *The Moral Point of View: A Rational Basis for Ethics* (Ithaca, NY: Cornell University Press, 1958), 181. See also his *The Rational and the Moral Order: The Social Roots of Reason and Morality* (LaSalle, IL: Open Court, 1995), 199.

12. P. F. Strawson, "Social Morality and Individual Ideal," *Philosophy* 36 (1961): 1–17, 11.

Actual Acceptance as Sufficient: A procedure *S* is eligible if *S* is accepted by all in the political community.

The "Exclusive View": Define the qualified public *P* as a subset of the actual political community; those in *P* possess the preferred account of the truth in political morality, *M*, such that *i* is a member of *P* only if *i* endorses *M*. Suppose that it is necessary and sufficient that *P* endorses procedure *S* for it to be eligible.

Estlund's "Qualified Acceptability Requirement": Specify the justificatory public *P* in some other way than by adherence to the true view of political morality, with the assumption that this specification is sufficiently wide that there are some "points of view" included in *P* which do not endorse *M*.¹³ Suppose that it is necessary that *P* endorses procedure *S* for it to be eligible.

It looks as if Estlund sets out to defend his Qualified Acceptability Requirement against these three competitors, but as I read this discussion the argument for it is very modest. Estlund's mode of argument, he tells us, is to respond to two objections (40)—there is very little positive argument for his favored Qualified Acceptability Requirement, so almost everything turns on the replies to the objections.

B. *Actual Assent Views*

The first objection comes from the proponents of Actual Acceptance as a Necessity, which claims that all qualified acceptance views, including Estlund's, are too exclusive, since they do not require acceptance by all actual points of view. Estlund's response is that his Qualified Acceptance Requirement only gives a necessary condition for justification, so it is consistent with Actual Acceptance as a Necessity since it too also only advances a necessary condition (47). He does

13. Estlund advances other specifications in terms of a qualified "doctrine" (43) and a "qualified objection" (47). It is not obvious that these formulations and that given in the text in terms of "points of view" are all equivalent. Suppose we take the route of simply identifying qualified *doctrines*. So we might be tempted to say "*O* is a reasonable objection if and only if it is made on the basis of a qualified doctrine." There is then no independent idea of a qualified objection. But this can lead to crazy results. Suppose a person reasons terribly on the basis of a qualified doctrine and so objects to a political regime. Is this qualified objection? It is hard not to insist on some standards of good reasoning, but then what is qualified depends not just on the doctrine but on the quality of the objector's reasoning. Perhaps the idea of a "point of view" intended to combine these, but there is certainly evidence that people who generally have a reasonable point of view can advance some pretty crazy specific claims. Do we want to count these as "qualified" even if we know they are based on a common inferential error? I leave these complexities aside here. See Gaus, *Justificatory Liberalism*, 54ff.; also see Gerald Gaus, "Reason, Justification and Consensus: Why Democracy Can't Have It All," in Bohman and Rehg, *Deliberative Democracy*, 207–42.

not reject Actual Acceptance as a Necessity. Given this, it is important to remember that, although the rest of the book employs only the idea of justificatory publics understood as qualified, the results concerning democracy are simply *pro tanto* since Estlund does not deny that additional defeaters can be forthcoming. He has not provided a case against the more inclusive justificatory public advocated by Actual Acceptance as a Necessity: he has merely insisted that his Qualified Acceptability Requirement is consistent with it. But, of course, according to Actual Acceptance as a Necessity, democracy may fail to be qualified since some member of the actual public may agree, for instance, with the Fifth Monarchy Men in the English Civil War, who interpreted Daniel's dream (Daniel 7) as indicating that there would be five great legitimate monarchies, the last of which would be that of Christ. They believed that the fourth monarchy, the Roman Empire, had been overturned by the Church of Rome, and so they were awaiting the fifth monarchy, the reign of Christ. Consequently, on the basis of their reading of the Bible, they denied the legitimacy of any further states and so of democratic states.

So much for the objection based on Actual Acceptance as a Necessity. Estlund must reject Actual Acceptance as Sufficient, since it is inconsistent with his Qualified Acceptability Requirement—it denies that qualified acceptance is necessary¹⁴—but there is little argument against it. Estlund holds that Actual Acceptance as a Necessity, which makes actual acceptance a necessary condition for legitimacy, although it is “radical and skeptical,” may still be true (46). Apparently Actual Acceptance as Sufficient cannot be true. If Actual Acceptance as a Necessity can be “true” but Actual Acceptance as Sufficient cannot be, there must be some damning case against Actual Acceptance as Sufficient that does not apply to Actual Acceptance as a Necessity. The odd thing is that Actual Acceptance as Sufficient seems more, not less, plausible (which is not to say that either is very plausible). Actual Acceptance as Sufficient at least minimizes the radical and skeptical tendencies of the actual acceptance view to which Estlund has drawn our attention since it does not require consent nor does it allow that, if there is consent, still other conditions might be necessary to show legitimacy. Actual Acceptance as a Necessity, in con-

14. Estlund interprets Actual Acceptance as Sufficient in a stronger way: “If the actual acceptance view were interpreted as both a necessary and a sufficient condition for legitimacy, then it would conflict with the qualified acceptability requirement” (*Democratic Authority*, 47). But there is no reason to interpret it that strongly; simply being sufficient is enough to generate the conflict with his necessary condition. In any event, if there is something deeply objectionable about actual acceptance being both necessary and sufficient, and this does not follow from the mere necessity requirement, it is a good hunch that the deep objection is to the sufficiency claim.

trast, opens the door to a very skeptical and radical theory: getting universal actual acceptance is required, but it may only be merely the beginning of our task! Once we accept that it may be true that any actual view has the normative power to defeat a justification, I don't see how we can reject as implausible the idea that universal actual acceptance is normatively sufficient. In fact, it seems far more plausible to hold that universal actual acceptance is sufficient than that it is necessary. If everyone actually accepted a political regime, it would seem somewhat odd and destructive to insist that it is not justified because there is some nonactual but possible good objection to it. (Why not let justificatory sleeping dogs lie and wait until one actually wakes up?) On the other hand, when we have a good normative justification for a regime that some actual people reject for bad reasons (think of our Fifth Monarchy Men), accepting their objections as possible defeaters is much less compelling. If Actual Acceptance as a Necessity is possibly "true," then so is Actual Acceptance as Sufficient. But then we still have in contention a principle that conflicts with Estlund's Qualified Acceptability Requirement.

C. Does the True Objection View Commit a Fallacy?

What, then, is the reason to reject our third contender, the Exclusive View, which defines the relevant justificatory public in terms of endorsing a true political doctrine? The proponent of the Exclusive View is apt to be a defender of what Estlund calls "epistocracy," which is his foil throughout much of the book. The defender of epistocracy holds that political decisions are legitimate only if they are made by the wise or the competent, or at least that the objections of the wise have more weight than those of the foolish. John Stuart Mill is often considered to be something of an epistocrat: "No one but a fool, and only a fool of a peculiar description, feels offended by the acknowledgment that there are others whose opinion . . . is entitled to greater consideration than his."¹⁵ The epistocrat may accept the idea that legitimate political procedures must not be open to qualified objection but then claim that the only qualified objections are those advanced by the wise. For purposes of simplicity, let us suppose we have a straightforward sort of epistocrat, one who holds that qualified objections can only come from those who have the correct view of the moral truth.

15. Mill, *Considerations on Representative Government*, 474. The relation of Estlund's account to Mill's is complex. Although Mill is in some ways the quintessential epistocrat, in other respects Estlund's view allows more elitism than Mill's. Although Mill advocates plural votes (the educated get more votes than the less educated), and Estlund criticizes this, Estlund's defense of "democracy" does not in principle seem to require universal voting rights, something that Mill endorsed for those people capable of representative government. See Estlund, *Democratic Authority*, 182, 209–11.

So this epistocrat may rely on the Exclusive View and argue that Estlund's Qualified Acceptability Requirement makes the relevant justificatory public too large and so it can disqualify procedures on the basis of objections from points of view that are not true (4).

Estlund's main argument against the Exclusive View seems to be that there is an argument for it that, he says, is "fallacious": "If it is the truth that matters, then it might seem that objections based on false views should not be allowed to defeat political justifications. Call this the *true objection view*: the only qualified objections should be true ones. *This line of reasoning is, however, fallacious.* If you love the truth, then you want to know what account of legitimate coercion is true" (51, second emphasis added).

So what precisely is the "fallacy"? It clearly involves some problem with (1) asserting that one who loves the truth will reject the Qualified Acceptability Requirement because it gets in the way of truth combined with (2) the possibility that the Qualified Acceptability Requirement is itself a truth. Estlund cannot be simply saying that proponents of the Exclusive View and of the Qualified Acceptability Requirement disagree about the truth of the latter, with the proponent of the Qualified Acceptability Requirement seeing it as true. If it were merely the fact that the true objection view had a false premise in its case, there would be no fallacy involved, simply a false conclusion. And in any event that would simply reiterate that the Exclusive View and the Qualified Acceptability Requirement are not consistent. Is Estlund claiming that one can consistently love the truth and accept his Qualified Acceptability Requirement (as true) and that the "fallacy" involved in the true objection view is a sort of (false) repressed premise that these are necessarily incompatible? Here, at any rate, is a form of the true objection argument that does not involve any repressed premises:

1. My love of the truth requires that I act to maximize relevant true beliefs and follow the directives of the reasons for action implied by moral truths.
2. The Qualified Acceptability Requirement is not true.
3. The Qualified Acceptability Requirement allows those with false beliefs to interfere with the goals in step 1.
4. I have decisive reason to reject any not true principle that interferes with the goals of step 1.
5. Therefore I have decisive reason to reject the Qualified Acceptability Requirement (or it is not a reasonable principle for me to endorse).

Now we can see a possible fallacy. The above argument appears to beg the question against the Qualified Acceptability Requirement by assuming it is false in step 2. The argument looks as if it supposes

that the Qualified Acceptability Requirement is false, but that, we may think, is precisely what the reasoning is supposed to show. One cannot give as a reason for rejecting a view as false an argument that has as a premise that it is false.

But no fallacy is actually involved in the above argumentative form; it does not beg the question. That is because the idea of “truth” (and what it means to “love it”) can mean different things, and so there is great room for ambiguity and equivocation.¹⁶ We come back to the problem of the scope of the concept of truth (Sec. II.B); Estlund’s minimal conception admits more types of claims as true than the objector may, and so they may disagree as to whether there is a truth about legitimacy. When a person loves the truth it will often be because she has a specific account of what it is (we might say that the concept of truth is part of her comprehensive doctrine). Recall the realist identified in Section II.B: to love the truth is to maximize knowledge of, and political fidelity to, a realm of nonnatural normative facts. Suppose Betty’s concept of truth—the thing she loves—is based on correspondence to such a realm. Now she may plausibly insist that the dispute about whether the Qualified Acceptability Requirement is politically justified—whether we have reasons to endorse it as a reasonable political principle—simply is not a dispute about truth: it is a dispute about political justification, warranted assertion, or some such thing. There is no reason why she cannot claim that the question of whether “the qualified acceptability requirement [correctly] defines legitimacy” —whether it is a justified conception of legitimacy—just is not about the correspondence of this claim to the nonnatural normative realm, and so on her view the Qualified Acceptability Requirement is not the sort of thing that could be *true* in the proper sense.¹⁷ Thus she denies that the

16. I believe the idea of “loving the truth” really is important to the argument—at least to its rhetorical force. (See also pages 4–5 where the idea is invoked three times.) The idea, though, is never explicated, and it certainly is not pellucid. On page 52 of *Democratic Authority*, Estlund seems to allow this: “Let us admit that the true objection view values truth in certain ways that a qualified objection view does not.” There is no further explanation what this might mean, but it points, I believe, to the critical difference on the basic conception of truth (Sec. II.B).

17. As we see, the objector might be tempted to use italics—the principle isn’t *true*—to make her point and to distinguish valid from true claims. Estlund employs such italics on page 59 of *Democratic Authority*, where he is considering a legal context in which (by stipulation) it is accepted that the defendant was far from the crime scene. A certain sort of contextualist may claim that as far as this legal context is concerned the defendant *was* far from the crime scene, so on the minimal theory of truth in this context that the defendant was far from the crime scene simply is true. To avoid this unwanted implication of his commitment to the minimal theory, I suspect, Estlund employs italics to signal a commitment to a noncontextual, more “objective,” notion of truth. As far as I can see, this indicates that Estlund’s argument (at this point in the text) may suppose a restriction

dispute about the Qualified Acceptability Requirement is a dispute about truth; and if that is not a dispute about truth, she cannot be begging the question when she claims the principle is “not true” in step 2.¹⁸ It is not a truth, but for her the question of whether it is justified (or authorized) is not the same as whether it is true. And she can entirely coherently say that those who care about the truth cannot honor a not true principle that blocks the pursuit of truth. She may intelligibly claim that to form our polity on legitimacy so conceived would greatly set back her great love—having a state that is maximally shaped by correspondence to such a realm. She may accept that it is an open question whether the Qualified Acceptability Requirement is “true” on the minimal conception, but this is not a conception of truth she employs, and on the correct theory of truth it is not the sort of principle that may be called “true” and it undermines the pursuit of the truth, properly understood. It constitutes a constraint on what truths the political order can respond to, the following of which, she argues, will sacrifice our attainment of a political order that, say, corresponds to many normative (real) truths. Those who do not love the truth as she knows it will get to block the polity's effort to conform to the demands of normative truth.

In short, Betty may be a type of political liberal: she holds that the political justification simply is not about truth, and she has strong views about the relation of acceptable principles to normative truth.¹⁹ Betty need not be committed to the minimal conception of truth; that is just one more controversial philosophical or comprehensive view. For Betty, as for many in moral and political philosophy, to call something a “truth” implies a certain ontological status or backing, which many justified claims or statements do not possess. Estlund's claim that a fallacy is being committed requires that the objector ac-

in scope that excludes contextualism—he certainly goes out of his way to show that the contextually correct “He was not at the crime scene” does not imply that it is *true* he was not at the crime scene (in this context).

18. This alerts us to hidden complexities in analyzing moral discourse. Once we admit the distinction between moral propositions that correspond to moral facts and purely conventional moral statements, we cannot infer from “not (x is F)” that “ x is F is false” if this latter claim is interpreted as “ x is F purports to correspond to a moral realm but fails to do so.” Relying on simple disquotational accounts of truth combined with two-valued logics can lead to apparent paradoxes. See, e.g., Ronald Dworkin, “Objectivity and Truth: You'd Better Believe It,” *Philosophy and Public Affairs* 25 (1996): 87–139. For an insightful analysis of these problems, see Kyle Swan, “Emotivism and Deflationary Truth,” *Pacific Philosophical Quarterly* 83 (2002): 27–81.

19. I have focused on this point as I believe it shows a deep flaw in Estlund's well-known argument that the political liberal is committed to seeing at least some political principles as true. See *Democratic Authority*, 53ff. For a fuller account, see David Estlund, “The Insularity of the Reasonable: Why Political Liberalism Must Admit the Truth,” *Ethics* 108 (1998): 252–75, 260ff.

cept his minimal conception of truth: she must accept a debate about whether to endorse a principle is ipso facto a dispute about the truth. But it is just that which Betty (and the political liberal) denies. So rather than showing that his account works with almost any concept of truth, we now see that Estlund requires that the objector accept his minimalist account of truth, which we have seen (Sec. II.B) may allow many more propositions as true than on the objector's preferred theory (which is what is happening in this case). This is, I think, important, for we see that the limitation in scope we noted in Section II.B is necessary for the argument against the main foil, the Exclusive View.²⁰ Many of those who care about the truth will reject this argument, with good reason.

In any event, Estlund concludes this discussion on a surprisingly modest note:

Nothing I have said shows that the qualified acceptability requirement, rather than the exclusive view, is true. My aim is only to point out that it, too, would be a truth. The exclusive view is not entitled to any advantage derived from the idea of loving the truth. Without that advantage the question becomes what basis there is for thinking that people are permitted to treat each other in that way: to coercively enforce laws even when one's only basis for doing so concerns matters about which people can reasonably disagree. This is a possible view, but it is not obviously true. (52)

How are we to understand the reference to "truth" here? I can only assume that we must understand it in the minimal sense. So to show that x is F is true is just to show that one may correctly assert " x is F." Now for Estlund to say that he has not shown that the Qualified Acceptability Requirement is true (rather than the Exclusive View) is just to say that he has not shown that a procedure is legitimate only if it meets the Qualified Acceptability Requirement rather than the Exclusive View. So at the close of the core discussion of chapter 3 Estlund explicitly acknowledges that he has not shown that the Qualified Acceptability Requirement is correct or that the Exclusive View is not. Thus one's puzzlement later in the book when he says: "Correctness theories cannot meet the qualified acceptability requirement. I take this to be conclusive against them" (99). Or, even more puzzlingly: "As I have argued in chapter 3, we should accept a constraint on political justification that prevents us from recommending simply

20. At one point Estlund claims that even if he allows that his criterion is not a truth, much of his argument will remain intact (58). I cannot explore this claim here, but it is certainly true that the argument claiming that the advocate of the Exclusive View is committing a fallacy fails without the truth claim, and it is the main argument against the Exclusive View.

the epistemically best method of decision, whatever it might be: political justification is specious if it relies on doctrines that could not generally be accepted by people with a wide variety of reasonable worldviews that will flourish under free conditions" (97).

We have just seen, however, that at the close of chapter 3 Estlund acknowledges that he has not shown that the Qualified Acceptability Requirement is correct or that the Exclusive View is not. So he has only "argued" for the Qualified Acceptability Requirement in an exceedingly modest sense. On my reading, we should treat it as an assumption, or a stipulation, on which the later analysis builds.

IV. THE QUESTION OF THE TEST, PART 2: THE EPISTEMIC TEST

A. The Increased Importance of the Epistemic Test

The upshot of our analysis of the first stage in answering the Question of the Test—justifying the principle by which we sort eligible from ineligible procedures—has not been encouraging. There is little argument for the Qualified Acceptability Requirement over the Exclusive View, and the main argument that is presented (that the "true objection view" is fallacious) depends on a specification of the scope of the concept of truth (Sec. II.B) that the objector almost surely does not accept. Now in itself this is not a decisive problem for Estlund's Qualified Epistemic Claim, that is, that "democracy . . . [is] [1] the best epistemic strategy [2] from among those that are defensible in terms that are generally acceptable" (Sec. I). If there is an effective demonstration of the first claim—that democracy has strong epistemic properties—then we may not be too concerned that the argument for the Qualified Acceptability Requirement (i.e., clause 2) is so tentative. In the terms in which I have framed the analysis, the answer to the first part of the Question of Test (the acceptability test) will not be all that important if democracy's passing the epistemic test shows it to have clear, interesting, epistemic virtues. On the other hand, if it should turn out that the only way democracy can be shown to be the best epistemic strategy from among those that are defensible in terms that are generally acceptable is that hardly anything but democracy passes the acceptability test (and so the acceptability test is doing most of the work in the argument), then the inconclusiveness of the argument for the Qualified Acceptability Requirement will undermine the case for Estlund's Qualified Epistemic Claim. Let us, then, consider the nature of the epistemic test.

B. Is the Epistemic Test Subject to Acceptability?

Suppose we have settled on some characterization of the qualified public (call this *P*) where membership in *P* does not depend on hold-

ing the correct or true moral view (we are finally putting aside the Exclusive View). And suppose we take *P*'s judgment as definitive in determining the eligibility of political procedures or regimes.²¹ The Qualified Epistemic Claim holds that democracy is (1) the epistemically best procedure (2) among those acceptable to *P*. Throughout *Democratic Authority* Estlund insists that a minimal requirement for a procedure or regime to be epistemically the best is that its answers must be better than random (e.g., 156, 160). Let us focus on this simpler, noncomparative claim about democracy's benefits; rather than showing that it is best with the set of eligible procedures, the main claim for democracy as *Democratic Authority* evolves is the non-comparative, better-than-random claim. Fundamental to understanding the Qualified Epistemic Claim is to see when a procedure produces better-than-random decisions.

But who must see this? The Qualified Epistemic Claim is ambiguous between a weak and a strong reading. On the weak reading, only clause 2 must not be subject to a qualified objection by any member of *P*. That is, on the weak reading democracy must be within the set of procedures to which there is no qualified objection, but then claim 1, the epistemic claim, is simply a matter of fact, in the sense of a proposition about democracy that Estlund maintains should be believed but is not subject to the qualified objection test. On the strong interpretation, not only claim 2 but also claim 1—the claim that democracy is “better than random”—must not be subject to qualified objection by any member of *P*. I think it is fairly clear that Estlund endorses the strong reading; as he says, “democratic authority rests on democracy's tendency to make better than random decisions, and better than alternative arrangements, *so far as can be determined within public reason*” (160, emphasis added). On the strong reading there must either be (1) some epistemic test *T* to which no member of *P* has a qualified objection and all members of *P* hold that democracy passes *T*, or (2) members of *P* employ different tests but all, employing their own tests, converge in their judgments that democracy passes.

I cannot see that we have reason to think that either option 1 or option 2 is the case. If that is true, on plausible conceptions of the public, there will be no agreement on the epistemic test or that democracy passes it. Once we examine the proposed test, as far as I can see, we must conclude that under almost any plausible account of the qualified public, the public will disagree about the test and whether democracy passes it. Of course there is a way to ensure that any doc-

21. I could not tell whether the focus is on procedures, as it seems in most places, or on regimes (as on 161, 167).

trine can be established within public reason: simply identify the qualified public as those who accept the test. But then, once again, the first part of the Question of Test is doing all the work, and we have seen that the results of that argument are rather inconclusive. I take it that a plausible defense of the Qualified Epistemic Claim must characterize P in a way that does not beg the question of the acceptability of the epistemic test in that manner. If P is simply defined as those who agree that democracy is better than random, the Qualified Epistemic Claim is vacuous.

C. Sensitivity and Discrimination

Suppose β is a procedure being employed to determine whether x is the case. As is well known, β 's accuracy has at least two dimensions: (1) If β comes up not- x when x is the case, it gives us a false negative. (2) If β comes up with x when not- x is the case, it comes up with a false positive. Estlund says that a procedure that avoids false negatives is characterized by *sensitivity*; one that avoids false positives has *discrimination* (113). Of course, it is easy to maximize accuracy along just one dimension: a test that always gives x as the answer will never give false negatives, and one that always says "not- x " will minimize false positives. The problem is maximizing over both dimensions.

In thinking about what counts as better-than-random performance, Estlund asks us to do the following:

Think about a basement detector, such as an electronic monitor for water in the basement. We can define random accuracy in the following way: the probability of water, given that the detector says water [p_w/Dw], is no higher or lower than the probability of water overall (its unconditional probability [p_w]). That would be a random level of discrimination. And the probability that the detector says water given that there is water [$p(Dw)/W$] is no higher than the probability that the detector says water overall [$p(Dw)$]—a random level of sensitivity. The conjunction of those two features defines random accuracy of the detector. (114)

So:

β randomly discriminates (i.e., randomly avoids false positives) if and only if $p_w/Dw = p_w$.

β is randomly sensitive (i.e., randomly avoids false negatives) if and only if $p(Dw)/W \leq p(Dw)$.

If the probability that there is water conditional on the detector indicating water is either greater or less than the probability that there is water, the detector nonrandomly avoids false positives. The idea of "nonrandomly" avoiding false positives should not be confused

with reliably doing so. If $p_{w/Dw} > p_w$, then the detector has some (nonrandom) reliability in avoiding false positives: it has some tendency, greater than the background probability that there is water, to indicate water only if there is water, and so it avoids indicating that there is water when there is no water. However, if $p_{w/Dw} < p_w$, the detector turns up more “false water” readings than would be achieved simply by using the background (i.e., unconditional) probability that there is water; this would be nonrandom but antireliable. If, as it were, we reversed the indicator lights, it would then have some reliability. Notice that it is an implication of Estlund’s analysis that, if $p_w = 1$ (there is always water), a perfectly reliable device (in the sense that it registers “water” if and only if there is water) is random on discrimination, since $p_w = 1$, and $p_{w/Dw} = 1$, so $p_{w/Dw} = p_w$. Estlund’s (initial) definition of random sensitivity is not parallel. If the probability that the device registers water when there is water is not greater than the unconditional probability that the device registers water, then the device has random sensitivity. Thus, if $p(Dw)/W < p(Dw)$, and so the device has less tendency to register water when there is water than its background probability of registering water (and so that when there is water it has less of a tendency to indicate water than it “usually” does), Estlund still deems it random, though in fact it seems antireliable (again, switching the indicator lights would render it somewhat reliable). This is puzzling, especially given that when Estlund turns to the definition of sensitivity in relation to political procedures (as opposed to devices in general), he abandons, without any explanation, this formulation of random sensitivity for one that parallels the definition of random discrimination, that is, “the probability, given that legislating x is a requirement of justice, that the procedure legislates x is no different from the unconditional probability that the procedure legislates x [i.e., $p(Dw)/W = p(Dw)$]” (115).

D. Views about Justice and Disagreements about the Relation of the Dimensions of Accuracy

Turning, then, to procedures about justice, Estlund makes accuracy about justice a four-dimensional concept by distinguishing false positives and false negatives on requirement questions from permission questions. Thus we get:

Requirement Sensitivity: Given that legislating x is a requirement of justice, it is very likely that the system legislates x . (114)

Requirement Discrimination: Given that the system legislates x , it is very likely that legislating x is a requirement of justice. (114)

Permission Sensitivity: Given that legislating x is permitted by justice, it is very likely that the procedure legislates x . (115)

Permission Discrimination: Given that the procedure legislates x , it is very likely that legislating x is permitted by justice. (115)

Estlund concludes that Requirement Sensitivity and Permission Discrimination are the most important desiderata. He argues that Requirement Discrimination is not of great importance because “it does not seem highly important that nothing be legislated unless doing so is required by justice. Some optional things might be positively good to do. On the other hand, there are some natural libertarian qualms about the possibility of too much legislation” (114). He argues against Permission Sensitivity on the grounds that “we certainly do not want every permissible law” (115).

As I noted at the outset (Sec. II.A), although we began with claims about moral truth, we are now concerned only with claims about justice. Early on it appeared that a utilitarian, upholding a strongly cognitivist conception of moral truth, would be within the ambit of the Qualified Epistemic Claim. But a utilitarian who holds that x should be legislated if and only if x maximizes aggregate welfare would have no reason to accept this analysis of what constitutes a better-than-random result. More generally, no theorist who characterizes the right in terms of maximizing good consequences is apt to endorse this analysis. And not even all deontologists will. The libertarian (whom Estlund mentions) does not simply worry about “too much” legislation but often endorses two theses:

1. Legislation is just if and only if it protects basic rights.
2. The state should do, and only do, what is required by justice.

Given theses 1 and 2, Permission Sensitivity collapses into Requirement Sensitivity and Permission Discrimination collapses into Requirement Discrimination. Unless libertarians and those having related natural rights views are defined out of the qualified public, some members of the public will hold that the requirement and permission dimensions are identical, while others will follow Estlund in excluding Requirement Discrimination and Permission Sensitivity.

Now perhaps these disagreements about what dimensions of accuracy are relevant will not make any difference to members of P 's evaluation of candidate procedures, but I see no reason to assume this. Someone who thinks that every act of legislation that is not required by justice is a false positive to be avoided certainly looks to have a very different view of an accurate procedure than someone who thinks that a large variety of legislation may not be required by

justice but is not against justice. The former will see all non-justice-required legislation as an act of injustice, and so “overlegislation” will be inherently unjust; she will see overlegislation as bad as underlegislation. In contrast, someone who follows Estlund will see underlegislation (in the sense of not doing what is required by justice) as much more serious than a tendency to overlegislate so long as this overlegislation does not violate some specific demand of justice. It would be a nice result to show that despite this deep disagreement on how the dimensions are related, qualified members of the public will converge in their judgment of a better-than-random procedure. I see no reason, however, to think this is so.

E. The Indeterminacy of the Very Idea of Requirement Sensitivity

Thus far I have been assuming that the qualified members of the public agree on the various dimensions of accuracy and what they require but that they simply disagree on how they are related. However, it seems most doubtful that qualified members of the public could agree on anything like a common interpretation of a test for any of the definitions of nonrandomness. Consider, for example, Random Requirement Sensitivity as it is applied to the political domain.

Random Requirement Sensitivity. [1] The probability, given that legislating x is a requirement of justice, that the procedure $[\beta]$ legislates x is no different from [2] the unconditional probability that the procedure $[\beta]$ legislates x . (115)

The problem is straightforward: qualified members of the public must agree on some way to estimate for a candidate procedure, β , both $p(\beta \text{ legislates } x)$ when x is just and $p(\beta \text{ legislates } x)$ unconditionally (or else, despite their disagreements, their judgments will converge). If Requirement Sensitivity is to be a test within public reason, and not just a formal definition with the key variables left undefined, members of P must have some common basis for reasoning about the probability of β 's “legislation indicator going on” conditional upon the matter being about justice and they must also agree on β 's unconditional probability of legislating. Surprisingly, Estlund provides no clue at all how they might do so. Let's consider some options.

Option 1: We might work via probabilities of members of different sets being legislated. Suppose we take a set of proposals $\{p_1 \dots p_n\}$ and then distinguish the set of justice requirements $\{j_1 \dots j_n\}$.²² If the probability that β will select some random element from the latter set (the conditional probability) is greater than the

22. Alternatively, we might define the sets inclusively, i.e., $\{j_1 \dots j_n\}$ is a subset of $\{p_1 \dots p_n\}$.

probability that it will select some such element from $\{p_1 \dots p_n\}$ (the unconditional probability), β is a better-than-random procedure for tracking justice. But a definition is not a test: qualified members of the public either must (1) converge on the same characterizations of $\{p_1 \dots p_n\}$ and $\{j_1 \dots j_n\}$, or else (2) their differences in their identification of the sets must not matter, and so they still converge in their conclusions. It is hard to see how either could be the case. Suppose Alf holds a view about justice that is not widely shared: he is a libertarian or a radical feminist whose set $\{j_1 \dots j_n\}$ of views is not widely shared by others, and he thinks that $\{p_1 \dots p_n\}$ is the set of proposals that are apt to be raised in a democracy because they stand some chance of being accepted. He will certainly conclude that the probability that democracy chooses from $\{j_1 \dots j_n\}$ is lower than the probability that it will select from $\{p_1 \dots p_n\}$: thus it is worse than random. And of course his rejection of democracy will be even more obvious if he thinks that every selection from $\{p_1 \dots p_n\}$ that is not a member of $\{j_1 \dots j_n\}$ is itself an act of injustice. In contrast, if Betty holds a widely shared conception of justice, she may be confident that the probability that democracy will select from $\{j_1 \dots j_n\}$ is higher than the unconditional probability that it will select from the random set of proposals $\{p_1 \dots p_n\}$.

Option 2: Of course the more we agree on the demands of justice, the more our judgments about better than randomness will converge. If, then, Estlund is willing to restrict the scope of the substantive views of moral truth for which the Qualified Epistemic Claim applies (Sec. II.A) by specifying a specific set of substantive conceptions that converge about some features of substantive justice, it looks like agreement on the test for better than randomness might be achieved. That is, they agree on $\{j_1 \dots j_n\}$. And, as we will see presently (Sec. IV.E), to some extent he does this. But that, alas, only solves half the difficulty: the identification of $\{p_1 \dots p_n\}$ remains a problem. Qualified members of the public must agree on the contours of that set before they can determine the unconditional probability that democracy will legislate (or else this disagreement must not matter). I see no compelling way to specify the set $\{p_1 \dots p_n\}$, which would allow us to identify the unconditional probability that any candidate procedure β will legislate. Some may propose that it should be all possible laws, but not only is it very hard to know what this set might look like, probably any system would come out as nonrandom if the unconditioned comparison set $\{p_1 \dots p_n\}$ is extremely large and includes many wacky proposals that no one would think of actually implementing ("Everyone wear pink underwear on Tuesdays!"). Suppose we get more realistic and select as $\{p_1 \dots p_n\}$ the set of bills that are apt to be introduced into the legislature. The idea then is that we take (a) the

probability that any randomly selected proposed law will be adopted and compare it to (b) the likelihood that a randomly selected proposal about justice will be adopted. If $b > a$, the procedure is sensitive to justice. But this determination will be greatly affected by the sorts of proposals that make it into the legislature. Compare:

The Restrictive System: Only proposals that have a significant chance of promoting economic efficiency or advancing justice make it into the legislature; proposals to advance efficiency are somewhat more likely to be passed than are proposals about justice.

The Loose System: All sorts of proposals make it into the legislature. It still is the case that economic efficiency legislation is more likely to be passed than justice bills, but both are more likely to be passed than the average somewhat nutty proposal.

If we define $\{p_1 \dots p_n\}$ in terms of proposals that are advanced, the Restrictive System will come out as insensitive to justice, the Loose System as sensitive, though it is possible that the two have identical legislation about justice.²³

Option 3: There are innumerable ways to specify the “unconditional probability” that democracy will legislate (i.e., the unconditional probability that β ’s “legislation indicator light will flash”), and it will matter which we select. Because trying to determine the correct unconditional set of proposals seems hopelessly complicated, it might be best to take another approach and seek to show that procedure β has some feature that makes it sensitive to the justice property (j) of a proposal x . Thus, when Estlund says that, according to the Random Requirement Sensitivity, “the probability, given that legislating x is a requirement of justice, that the procedure $[\beta]$ legislates x is no different from the unconditional probability that the procedure $[\beta]$ legislates x ,” we might interpret this as saying that a procedure is random in this regard when it is equally likely to legislate when some specific proposal x does and does not have property j . The idea could be put in terms of possible worlds. Imagine world ψ , in which x is a matter of justice, and take the nearest possible world ω , in which there is a counterpart x but it is not a matter of justice; if the probability of β legislating x is greater in ψ than in ω , then β meets Requirement Sensitivity. This avoids the problem of finding a suitable comparison

23. It is worth noting that on this interpretation a utopian classical liberal political system in which only just proposals were made, and in which the legislature passed all such proposals, would be deemed to have random sensitivity. The $p(\beta$ legislates from $\{j_1 \dots j_n\})$ would equal the $p(\beta$ legislates from $\{p_1 \dots p_n\})$, since $\{j_1 \dots j_n\} = \{p_1 \dots p_n\}$. Because it would score as random on this dimension, it could be strongly dominated by less perfect systems which conceivably could score as better than random on all dimensions!

set of unconditional β outputs, but it raises its own host of problems: not only must members of P agree on the nearest possible world and what β would do in it, but the identification of x counterparts would often be highly controversial. Members of P would need to identify “ x ” when it has the property of being about justice and when it does not: since the property of being about justice will be an essential property of many proposals, we can predict intractable dispute whether a nonjustice counterpart can be identified.

All this is, admittedly, complicated stuff, but that is because the core epistemic test that Estlund proposes is a formal idea: a better-than-random result. It will not do, I believe, to commence an analysis with a strong claim about the importance of public acceptability of possible procedures (the first stage of the Question of Test) but then, at the next stage (the epistemic test), appeal to a test that can only be adequately spelled out employing highly contentious formal reasoning. Political philosophers are, perhaps, apt to take “Does democracy have a better-than-random chance of yielding the right answer?” as at least posing a clear question, but I have tried to show that the very meaning of the test is highly controversial (indeed, I think, hopelessly indeterminate) and that this matters for judgments about the epistemic performance of democracy.

F. Primary Bads: Severely Restricting the Scope of the Justice Set

I mentioned above (Sec. IV.E, option 2) that some of the problems in applying the key idea of the “better-than-randomness” test (and it is a key, for without it we cannot even interpret the Qualified Epistemic Claim) can be mitigated by narrowing the range of judgments about justice, $\{j_1 . . . j_n\}$. As I noted near the outset (Sec. II.A), Estlund does this, focusing on democracy’s ability to avoid what he calls the “primary bads”: “war, famine, economic collapse, political collapse, epidemic, and genocide” (163). It is perplexing that, after a focus on justice as the core evaluation by which better than randomness is determined, we now see that the crux of the shared judgments is about avoiding bad outcomes. To bridge the gap between justice and badness I would have thought we would need an argument that it is always an injustice to bring these bads about. But we know that it is not always an injustice to bring war about, for there are just wars, and I take it that not all just wars are wars of self-defense. A state can justly choose to start a war; it is unjust wars that certainly must be avoided. (Sweden avoided the primary bad of the Second World War; Norway did not. What does this tell us?) But then, as Estlund suggests at one point (268), we need an account of justice, not just bads;²⁴ but just

24. Or, as he says in another place, “unfortunate bads” (163).

war theory is by no means uncontroversial. So the list is more controversial than it looks. Moreover, we do not understand the full causes of many of these bad events, and avoiding them may involve complex strategic interactions with other actors beyond the power of any one polity to determine the outcome. It is certainly puzzling that we get no justification whatsoever for the omission of the traditional uncontroversial injustices that states can themselves avoid, such as arbitrary arrest and imprisonment of citizens, basic violations of the person (e.g., torture and rape by government officials), and so on.

Leaving these puzzles aside and returning to the test of randomness, Estlund's claim is that one who questions that at least "some democratic arrangements are epistemically better than random" is "immediately in an awkward position" because the claim made in favor of democracy is so "modest" (168). The modesty involved here is partly that Estlund has only taken a small subset of "justice" requirements. But we have seen (Sec. IV.D, option 2) that identifying even the entire set $\{j_1 . . . j_n\}$ is not sufficient to apply the better-than-randomness test; it is by no means obvious how we are to evaluate this "modest claim" without a good handle on the unconditioned probabilities. What is the unconditional baseline used to measure the performance of democracy? A problem in coming to grips with this matter is that the question has been changed: rather than talking about the probabilities that the legislature's "legislation light" will blink unconditionally and when there is a matter of justice, we are now considering when a state of affairs will occur. Because Estlund does not tell how to translate the test into this new context, we are pretty much on our own again.

Let us see if we can get clear about the claim about democracy and primary bads in a way that retains the key idea that democracy excels at avoiding injustice. Following the suggestions of the third option in Section IV.D, take as the unconditional baseline a world in which the legislature has no conception of justice; it only has the concept of good and bad. The probability that a democratic legislature will avoid the primary bads in this world gives the unconditional probability that a legislature will act to avoid these bads. We then compare this to a world in which the primary bads are seen as a matter of justice; the probability that the democratic legislature will avoid these primary bads would be the conditional probabilities. The most natural extension of the epistemic test would then be that a democratic legislature in the second world would do better than a democratic legislature in the first, thus showing that it has a better-than-random sensitivity to the injustice of the primary bads and it is this which allows it to avoid them in a better-than-random way. But

now, however, it is plausible to think that the unconditional and conditional probabilities will be the same. Just because these are such awful states of affairs the injustice of them may not make any difference. Indeed, one can argue that once we consider justice they may become more likely—a state that is concerned with justice may start a war to make the world safe for democracy, and this war may cause political collapse. Think of it this way: in the unconditional situation, we have only utilitarians who care nothing about justice but who care a lot about avoiding bad states of affairs; in the conditional situation, we have a legislature seeking to track justice deliberating about the primary bads. If democracy is obviously a better-than-random way of avoiding the primary bads because it tracks justice, it would have to be the case that the utilitarian legislature obviously will do worse. But concerning precisely the issues on which Estlund has based his modest case for the epistemic benefits of democracy—avoiding really bad states of affairs—this is dubious.

What if we entirely leave aside the idea that sensitivity to justice is essential to the Qualified Epistemic Claim and simply focus on avoiding these bad states of affairs? So the claim would be simply this: a polity with a democracy has a better-than-random chance of avoiding these states of affairs, period. But again, we cannot even interpret this without an unconditional comparison class. What might be the unconditional probability that these states of affairs will be avoided? Readers might be tempted to think that this is a simple problem: take the total number of states in the world (N) and work out the frequency of the primary bads in them (f_N), then take their frequency within democracies (D), and if $f_D < f_N$, then democracy is a better-than-random way of avoiding them. Unfortunately, this could only show that democracy is better than average, and it is still possible that all are worse than random. In any event, this way of thinking does not distinguish causal correlation from spurious correlations. It would also follow that countries with high iPhone usage are good at avoiding the primary bads; that hardly gives a reason to applaud the way iPhones allow a better-than-random avoidance of famine. If we begin to look seriously at causal factors, I would very much suspect that, once we distinguish their separate effects, effective markets and property rights are a much superior way of avoiding famine and disease than is democratic deliberation. Many economists would insist that deliberating about claims of justice would have no real effect on the outcome (some say it would make matters worse).²⁵ This is, of course,

25. See, e.g., Louis Kaplow and Steven Shavell, *Fairness versus Welfare* (Cambridge, MA: Harvard University Press, 2002).

disputable, but as far as I can see this basic case for democracy concerning primary bads is supposed to be, in some sense, beyond reasonable disagreement (169). It is not.

V. THE QUESTION OF MECHANISM

We have been considering the test for democracy's competency at tracking the truth, as well as some ideas about whether it might pass the test (whatever that test turns out to be). Now just as I argued (Sec. IV.A) that the modesty of the argument for the acceptability requirement would not undermine the Qualified Epistemic Claim (Sec. I) if the epistemic test could be clearly stated in a way that shows democracy passes it, the inability to specify the test would not lead us to lose faith in Estlund's epistemic case for democracy if we had a nice account of the mechanism by which democracy yields true belief.

It may appear that a plausible mechanism is provided by Estlund's imaginary deliberative model. It is this model which, in the end, is supposed to provide the core epistemic credentials of democracy and so on which the Qualified Epistemic Claim must ultimately rest. The democratic ideal is defined in terms of a model of deliberation which, we are told, "serves as a kind of template by which to mark and measure deviations and devise epistemically remedial responses" (175). According to the ideal model:

1. Everyone has full and equal access to the forum.
2. Everyone has the same chance to speak as everyone else.
3. People only say things that they believe will help others appreciate the reasons to hold one view or another among those that are in question.
4. Anyone whose interests are at stake in the decision is either present or represented by an effective spokesperson.
5. Everyone has as much time to speak as they wish.
6. Everyone has equal bargaining power.
7. Everyone equally credits and attends to the contributions of all others.
8. Everyone recognizes (or tends to recognize) a good reason when they see it.
9. Participants strive to address the "devil's advocate." (175–76)

Estlund claims that: "[1] a deliberation conducted under these conditions is likely to have a significant tendency to make decisions that are morally right by standards that are independent of the results of this deliberation. [2] Primary bads would be fairly reliably avoided; overall decisions would be better than if they had been made randomly" (176).

Now the ideal deliberative process has epistemic value only if it can be shown that deliberation (236) is what produces the better-than-random (whatever that turns out to mean) results rather than, say, an initial above-random competency of each person (which some interpretations of clause 8 readily suggest). That is, we require a demonstration of:

The Amplification Claim: Whatever the predeliberation average reason-tracking competency each individual has on her own, the average (or group) reason-tracking competency after some period of ideal deliberation is significantly higher.

As Estlund recognizes, the Condorcet Jury Theorem is so attractive to an epistemic theorist because it provides a precise account of democracy's "competence-amplifying effects" (234). However, Estlund acknowledges that he advances no "detailed account of when reasoning together will improve group competence" (234). To be sure, he provides a few examples of amplification and employs an analogy, but he commits to no systematic account of how it occurs even in the ideal deliberative model. This makes the evaluation of the Amplification Claim, and ultimately the Qualified Epistemic Claim itself, exceedingly difficult. My worry is that, without committing to any account, we do not know the range of claims for which amplification might occur, and so we do not even know when the ideal epistemic model plausibly produces amplification. But that is the heart of the argument for the Qualified Epistemic Claim.

To better see the problem, consider one candidate for the foundation of the Amplification Claim:

The Ideal of Convergence: As an ideal deliberator, Alf holds that if he believes p and Betty believes not- p , at least one of them is in error. As an ideal deliberator, Alf supposes that deliberation among ideal deliberators should tend toward convergence of belief; his and Betty's disagreement is a reason to reevaluate their p -related beliefs. Because, according to claim 8, they recognize good reasons, the result of this reevaluation is that they tend to converge on more accurate/better p beliefs.

If the Ideal of Convergence holds, then I think the Amplification Claim has some plausibility.²⁶ But once we see a possible basis of the

26. However, there are still formidable difficulties in its way. Unless individuals can process the information they get in large-scale deliberation, rather than amplification there may simply be informational overload. Think here of the F. A. Hayek's argument that individuals simply are unable to collect and process all the relevant information about how to achieve optimality among individual plans ("The Use of Knowledge in Society," *American Economic Review* 35 [1945]: 519–30). Even full recognition of the set of relevant reasons may do no good, since the computational process required to make sense of all

Amplification Claim we immediately also see some limits. Consider a problem with which we began: indexical truth claims (Sec. II.B). The Ideal of Convergence is dubious when true claims are cognitive but indexical. If, say, Alf holds that his p claim is indexical to his moral community, and Betty's not- p claim is indexical to her community, then they do not suppose that this disagreement needs to be investigated and, hopefully, resolved. On some views such indexical relativism is the only sound basis for "faultless disagreement" where there is no mistake by either party to be discovered and so no good reason for convergence.²⁷ Given this I see no compelling reason to accept the Amplification Claim for such indexicals. Only in a special case—when the deliberating group is coextensive with the indexed group—is it plausible that the Ideal of Convergence would hold given cognitive but indexical moral claims. That would require that political communities be coextensive with moral communities. This is not itself an objection to Estlund's account, since at least on some views indexicals do not bear disquotation and the minimal truth claim requires disquotation, but it does point us once again to the limitations on scope of admissible truth claims (Sec. II.B). If, supposing Estlund is committed to the Amplification Claim based on the Ideal of Convergence, then many cognitive theories with indexical moral truth claims (perhaps Baier's and mine) lie outside the ambit of the Qualified Epistemic Claim. If Estlund wishes to avoid this exclusion, then it is hard to see how he can base his case for the Amplification Claim on the Ideal of Convergence. But then what is amplification based on? More generally, we cannot know the range of claims for which the ideal deliberative model would induce amplification without having a clear idea of the mechanism producing amplification, and so we are left in the relative dark about why democracy has its supposed epistemic benefits.

this information may be beyond any central processing ability, and so deliberations about optimality may simply increase a person's store of false beliefs. In response to Hayek's point, Estlund criticizes overreliance on self-interest, markets, and "the epistemic invisible hand" (*Democratic Authority*, 177) and then asserts that "there is no devastating case here against the epistemic value of deciding questions politically rather than through markets" (*ibid.*, 179). But whether markets can do the job better is, I think, something of a red herring; perhaps they are both bad. The real question is whether Estlund has an adequate reply to Hayek's analysis of the computational complexities of large-scale optimality problems. On some views of the political truth, seeking the "better," or the common good, may well be an optimality problem. Hayek's argument is that deliberative processes in which individuals seek to identify the global optimum are hopelessly cognitively demanding. I do not believe these problems can be easily dismissed or simply "denied" (*ibid.*, 179) without a sophisticated analysis.

27. Kölbel, "Faultless Disagreement."

VI. CONCLUSION

I have tried to delve carefully into Estlund's case for the Qualified Epistemic Claim. It is a case he presents in a detailed way, and the details deserve to be examined. We have seen that the Question of Scope is answered differently at different points in the analysis; the scope is considerably restricted as the book proceeds, and it sometimes changes unpredictably (e.g., from moral truth to justice to bads). Neither the Question of Test nor the Question of Mechanism is given an adequate answer. Given (1) the inconclusive argument against the Exclusive View (and the way it seems to require that members of *P* accept restrictions on the scope of truth claims), as well as the very modest argument for a broader conception of the qualified public, (2) the grave difficulties in even interpreting the idea of a better-than-random legislative output about justice and its inability to function as a test within public reason, and (3) the admitted lack of a systematic analysis of the Amplification Claim even in the ideal deliberative model, I conclude that Estlund has presented some suggestive considerations in favor of the Qualified Epistemic Claim but not, I think, anything like a convincing case for it. The empirical evidence indicates that democracy does well at avoiding some very bad and unjust policies, and it is by far the best way humans have discovered to organize themselves politically without systematic and serious oppression. But I cannot discern in *Democratic Authority* a well-defended case in public reason for democracy as generally a better-than-random procedure for legislating according to political truth, whatever that may be. In the end, Estlund's intelligent and original work has failed to assuage my Millian worries.

But perhaps I simply ask for too much. In evaluating Estlund's case for the Qualified Epistemic Claim I wanted to see whether it assuaged my worries that Mill might be right. That is, I have searched for compelling reasons to believe that democracy really has interesting epistemic features. Is it too much to ask for such reasons from a philosophical treatise? My good friend Stanley Benn began his great book, *A Theory of Freedom*, by describing philosophy as about identifying reasons for "appraising beliefs, arguments, and proposals to act."²⁸ In the sciences, bold conjectures have an important place, for there are also data that help us adjudicate which of these bold conjectures we have reason to believe and which we do not. In political philosophy, as in all philosophy, bold conjectures cannot substitute for arguments providing reasons to believe the premises of an argument and its conclusions. Indeed, in philosophy, advancing argu-

28. S. I. Benn, *A Theory of Freedom* (Cambridge: Cambridge University Press, 1988), 1.

ments that provide as strong as possible reasons to believe is the main discipline on our discipline, checking our aspirations and flights of imagination by providing as near as we can get to a common tribunal of public reason. *Democratic Authority* is a book filled with insights, tantalizing ideas, and bold conjectures (when reading it I was reminded of Robert Novick's *Anarchy, State, and Utopia*). But at least on my old-fashioned and less playful conception of political philosophy, it falls short. If we care about the truth, we want reasons to believe it.