

COPYRIGHT NOTICE:

David M. Estlund: Democratic Authority

is published by Princeton University Press and copyrighted, © 2007, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

Follow links for Class Use and other Permissions. For more information send email to: permissions@pupress.princeton.edu

Democratic Authority

One of the two, as the wiser or better man, has a claim to superior weight: the difficulty is in ascertaining which of the two it is.

John Stuart Mill, *Considerations on Representative Government*

DEMOCRACY CAN SEEM TO EMPOWER THE MASSES without regard for the quality of the political decisions that will result. Concern for the quality of decisions can seem to lead in an antidemocratic direction, toward identifying and empowering those who know best. Partly for these reasons, philosophical treatments of democracy's value have often tried to explain why politics should be democratic even though democracy has no particular tendency to produce good decisions. I believe these accounts are weak, and I want to put democratic convictions on more secure footing. My goal is to show how a concern for the quality of political decisions, properly constrained by other principles, supports democratic political arrangements. Rousseau pleads, "All my ideas are consistent, but I cannot expound them all at once."¹ In this synoptic essay I try to present, all at once, the broad shape and main points of a more elaborate book-length argument. I hope it has some value as a self-contained essay, but it is also intended as a prelude to the longer argument of the book, making certain points and terms available for easy reference even before they are treated fully.

A FRAMEWORK

Before turning to democracy, I begin with the idea of a philosophical framework. Political philosophy, as with some areas of ethics, is easily distorted by the ever-present thought that it might be of practical importance. Practical applications of philosophical ideas require engagement with a lot of nonphilosophy, and the danger is not just that philosophers are not normally especially good at the relevant nonphilosophical

areas of inquiry. Even if they were, there are risks involved in trying to treat both kinds of questions in the same work. In the hurry to make a practical proposal it is easy to lose sight of the philosophical problems, and so to lose sight of whether and how they have been solved. Since even long-standing problems have, so often, not been solved (philosophy seems to be harder than science in this way), the idea that something is gained if political philosophers explain how to put their ideas into practice is hard to understand. If the more abstract arguments are of any value, then it would be good if someone takes up the further questions about what they might imply or recommend in practice, and I do not mean to denigrate that task. My main concern, however, is with the question whether certain points at a fairly abstract level are of value in the first place, whether they solve or at least contribute to the solution of important problems that lie at the more abstract level. As a result, few institutional specifics are offered here, and when they are they are mainly meant as illustrative examples, not as prescriptions.

There is a second aspect to the limitation I have in mind by providing only a philosophical framework: detailed factual information, while occasionally useful, is far from the center of our concerns. The focus of the argument is on very general questions of authority and legitimacy in a political community, terms that will be used in very specific ways. By *authority* I will mean the moral power of one agent (emphasizing especially the state) to morally require or forbid actions by others through commands. (To forbid x is to require not- x , and so I will usually simply speak of the moral power to require.) By *legitimacy* I will mean the moral permissibility of the state's issuing and enforcing its commands owing to the process by which they were produced. If the state's requiring you to pay taxes has no tendency to make you morally required to do so, then the state lacks authority in that case. And if the state puts you in jail for not paying, but it is morally wrong for it to do so, then it acts illegitimately. Even without authority or legitimacy, of course, the state might yet have enormous power. But we are not investigating brute power as such, since brute power is not a moral thing. Like a knife, it can be used rightly or wrongly. The moral questions about the use of knives are not much about the details of what knives are like, and the moral questions about the uses of power are not much about the exact nature of actual power.

Even without great emphasis on empirical studies about politics, or on the details of institutional design, the philosophical tasks alone lead

to a rather elaborate theory. In this essay we look only at the main points and lines of argument.

MAKING TRUTH SAFE FOR DEMOCRACY

The idea of democracy is not naturally plausible. The stakes of political decisions are high, and the ancient analogy is apt: in life-and-death medical decisions, what could be stupider than holding a vote? Most people do not know enough to make a wise medical decision, but a few people do, and it seems clear that the decisions should be made by those who know best.

While it makes good sense for us to defer to someone who we have reason to think is a medical expert, the doctor's right to make decisions and perform procedures on us comes mainly from our consent, not from the doctor's expertise. Politics is different, since most of us have never consented to the political authority of the government that rules over us. Authority would need some other basis, and expertise has long been a tempting one. It is important to see that authority does not simply follow from expertise. Even if we grant that there are better and worse political decisions (which I think we must), and that some people know better what should be done than others (we all think some are much worse than others), it simply does not follow from their expertise that they have authority over us, or that they ought to. This expert/boss fallacy is tempting, but someone's knowledge about what should be done leaves completely open what should be done about who is to rule. You might be correct, but what makes you boss? Perhaps this approach to political justification, which draws heavily on the views of John Rawls,² points in a democratic direction.

Consider religious points of view. One of the contending views might be correct, with all its implications about what should be done politically. Suppose your religious point of view is not the true one. If you think the knowers should rule—if expertise entails authority—then you must think that those with the true religious perspective, whoever they might be, should rule even over people like you who mistakenly doubt that they are the knowers. Their being experts (so to speak) certainly counts in favor of your *accepting* them as rulers, but as it happens, you do not realize they are experts. This doubt appears to block the inference from their expertise to their authority. The problem isn't exactly that you haven't consented, and we're not assuming that

consent would be required to legitimate rule. It's about what you believe: you do not believe that they are experts.

The problem that arises for this line of argument, however, is that there is not much that will be believed by everyone, and if legitimate authority depends on there being a justification accepted by everyone, it will be hard to find much legitimate authority. But this is not a plausible constraint anyway. Why should the objection of someone who is, say, crazy or vicious carry that much moral weight—enough to defeat a justification even if it is acceptable to everyone who is not crazy or vicious? Rather than say that a justification must be acceptable to everyone, we might try saying that it must be acceptable to everyone except the crazy, the vicious, and . . . well, there might be other things that it makes sense to put in here. We will finesse the details by simply saying that there will be some list of things that disqualify certain points of view for this purpose. That is, some points of view are such that objections that depend on those disqualified points of view are not capable of defeating a proposed political justification. People who believe that their own race has a right to rule other races, or who simply desire to subordinate other people to their power, will not accept certain principles about moral and political equality. But objections stemming from those unreasonable points of view are morally weightless.

By calling some views qualified and others disqualified, we tempt objectors to accuse us of being exclusive or elitist. This is a widespread confusion, in my view. Since all we've said so far is that not all points of view are qualified, a more inclusive view would have to say that all points of view are qualified. Justifications must simply be acceptable to everyone. This is an oddly amoral view, in which otherwise sensible lines of justification are unavailable if they are not acceptable to Nazis. If anyone believes this, we would need to hear a lot more in support of it.

The other argument against treating some grounds of objection as disqualified says that it is too inclusive—that it counts too many (not too few) lines of objection as qualified. In particular, some say that a point of view shouldn't defeat a proposed justification unless that point of view is true, whereas our approach allows that some views are qualified even though they are not true. This objection might seem to be the proper view for any lover of the truth. Some objections to proposed political justifications are correct and others are not, and lovers of truth should only care about the correct ones, so the story goes. But lovers of the truth want to know what the truth is about justification itself as well, and that requires that we determine whether only true points of

view or also some others are qualified to defeat proposed justifications. If the truth is that justifications must be acceptable to all qualified points of view, including many that are not true, then lovers of the truth should accept this view of justification. We cannot settle it here, but this approach would explain the thought that even if the pope has a pipeline to God's will, it does not follow that atheists may permissibly be coerced on the basis of justifications drawn from Catholic doctrine. Some non-Catholic views should count as qualified for this purpose even if they are mistaken. This itself asserts a truth about justification, as lovable a truth as any other.

FAIRNESS AND PROCEDURALISM

All this talk about truth will drive some readers crazy. Some will deny that there are any truths about what should be done politically, but few mean this in a way that would cause any difficulty for what I have said. The nature of truth is a fascinating philosophical matter, and truth in value judgments raises interesting questions of its own. But when I speak of moral truth here, I mean only the following very minimal thing: if gender discrimination is unjust, then it is *true* that gender discrimination is unjust. Not many readers will think nothing is just, unjust, right, wrong, and so forth, and so they accept that there are moral truths in the sense that concerns me.

Some will worry about whose view of the truth we are talking about. But we are not talking, initially, about anyone's view of the truth, but about the truth (whatever it might be). We are supposing that some things are unjust, some right, some things vicious, and so on, regardless of what anyone thinks about them. Then we say that some people have erroneous views about these matters, some other people less so. So far we are not endorsing any particular *view* of the truth. We have not said which things are true about these matters, or who might know the truths better.

This way of talking about truth makes it pretty hard to deny truth in political matters or to deny that some know it better than others. We cannot resist the move from expertise to authority by denying that there is expertise, then. I have proposed to avoid that move by denying that there is expertise that is generally acceptable in the right way even if it is genuine. But having acknowledged that there are truths about the high-stakes matters that are present in politics, we must ask whether its

discovery plays any role in the best account of how political authority and coercion would be justified. Is there an epistemic or truth-seeking dimension to the best account, or can we do without that?

I believe we cannot do without it, but there is a simple and influential approach that tries to. Why not understand democracy as a way of giving every (adult) person an equal chance to influence the outcome of the decision? The justification of the outcomes would be in terms of the familiar idea of the fairness of the procedure that produced the decision. That way we would not need to make any claims about the decision tending to be good or right or true. Democracy, after all, does seem like a fair way to make decisions, at least as an aspiration. People are given equal rights to express their political opinions, and equal rights to a vote. Should we say that the fairness of the procedure is the explanation of democracy's moral importance?

We have not said exactly what would make a procedure a fair one, but so far it looks like democracy is one fair procedure, and choosing between two proposals by flipping a coin is another one.³ If that is right, and if fairness is the main basis of democracy's importance, then why not flip a coin instead? It is much cheaper and easier in so many ways. We would not need to expend resources on campaigns, televised debates, public political discussion, or all the time and work involved in holding a vote. For example, we could let the slate of candidates or issues be determined in whatever way they otherwise would, except at any stage that there would normally be a vote we substitute a random selection process, even at the final stage that would normally be an election or a referendum. It is a perfectly fair procedure, at least if this means giving each person an equal chance at changing the outcome. All have an equal chance, since no one has any. If the value of democracy is its fairness, this random procedure should be just as good.

Of course, this is impossible to accept. There is something about democracy other than its fairness that contributes to our sense that it can justify authority and legal coercion. A coin-flipping procedure would not justify these, at least not to the same extent. One natural hypothesis about why we actually want people's views taken account of by the process is that we expect people's views to be intelligent—maybe not to any high standard, but better than a coin flip. This introduces an epistemic dimension, and it is important to see what this would and would not commit us to.

EPISTEMIC PROCEDURALISM

The biggest objection to bringing in the epistemic dimension is that it might tend to justify rule by the knowers—what we might call *epistocracy*. But I hope to have forestalled that worry by pointing out that even if there are knowers, it might well be that their status as knowers is not generally acceptable in the way that political justification would require. So now the question is how to bring in knowledge without privileging any class of knowers.

There is, however, the nagging thought that even if the knowers aren't generally acceptable, they do exist. And they might simply be a more accurate source of knowledge about what should be done than any democratic procedure could ever be. I don't want to deny this. The question is not how democracy might be the best epistemic device available, but how it might have some epistemic value in a way that could account for the degree of authority we think it should have. If you or someone whose opinions you trust is a knower, then the results of a modestly epistemic democratic procedure will not give you especially important epistemic reasons to believe the outcome is good or correct. But epistemic reasons are not what we need. The hope is to show how democracy yields moral reasons to obey the law and a moral permission to enforce it. We should not assume that there could only be such duties and permissions in cases where the procedure actually got the right answer. So we should not assume that authority and legitimacy lapse just whenever the procedure gets a wrong answer. That surely would not generate legitimacy and authority for the general run of democratically produced laws. Obviously, many of them are unjust or otherwise morally mistaken.

A good model for the structure I have in mind is a jury system. When it is properly done, a jury trial seems to produce a verdict with legal force, but also with some moral force. If the defendant is exonerated, then other people will have a moral duty not to carry out private punishments. If the defendant is convicted, then the duly appointed jailer will have a moral duty not to set him free. I assume that, at least within limits, these moral implications do not depend on the verdict being correct. If they did, then we should all ignore the verdict and use our own judgment about whether the defendant was guilty or not, and privately punish the truly guilty and open the cells of the innocent. Jailers who appoint themselves judge and jury, and vigilantes who appoint

themselves executioner, seem to be acting immorally when there has been a properly conducted (though always fallible) jury trial.

The jury trial would not have this moral force if it did not have its considerable epistemic virtues. The elaborate process of evidence, testimony, cross-examination, adversarial equality, and collective deliberation by a jury all contribute to the ability—certainly very imperfect—of trials to convict people only if they are guilty, and not to set too many criminals free. If it did not have this tendency, if it somehow randomly decided who is punished and who goes free, it is hard to see why vigilantes or jailers should pay it much heed. So its epistemic value is a crucial part of the story. Owing partly to its epistemic value, its decisions are (within limits) morally binding even when they are incorrect.

On this account, the bindingness and legitimacy of the decisions are not owed to the correctness of the decisions, but to the kind of procedure that produced them. Still, a central feature of the procedure in virtue of which it has this significance is its epistemic value. I call this theoretical structure *epistemic proceduralism*. This is just the structure I explore in the case of democratic procedures for making laws and policies generally. Democratically produced laws are legitimate and authoritative because they are produced by a procedure with a tendency to make correct decisions. It is not an infallible procedure, and there might even be more accurate procedures. But democracy is better than random and is epistemically the best among those that are generally acceptable in the way that political legitimacy requires. The authority and legitimacy of its laws often extends even to unjust laws, though there must be limits to this. The procedure does not give us great epistemic reasons for our opinions about justice. For that purpose we may each appeal to whatever sources and methods we think best, without the need for these to be generally acceptable.

Suppose the tax laws charge me more than is really just. But the laws were passed in a properly functioning democratic system with ample opportunity for discussion and debate. My objections were aired and answered, even though I think the answers were inadequate (and even suppose I'm right about this). The resulting laws charge me an unjust amount of tax. (Assume it does not take so much that I cannot still live a decent life. We will not try here to locate the limits, but there must be some.) According to epistemic proceduralism, the law is legitimate and binding on me even though it is unjust, and this is owed partly to the fact that the procedure has epistemic value that is publicly recognizable.

This gives something of the structure of epistemic proceduralism's account of democratic authority. We need to say more, though, about

why we should think that a procedure with these features does have authority, the moral power to require action.

DEMOCRATIC AUTHORITY WITHOUT CONSENT

A traditional view says that there is no authority without consent. The state is not in a position to lay obligations on me unless I voluntarily and knowingly agree to their having that moral power. The main weakness of this approach is that it does not seem to account for the state's authority over very many people, since most people never consent to the authority of their state. Locke argued that a person consents merely by enjoying the benefits of the state, but this seems to make hash out of the idea of consent. Since you could benefit from the state without realizing you were consenting, you could consent without realizing it. But I would have thought that the whole point of requiring consent is to let the person decide whether to be under the authority or not. If you could consent to authority without realizing it, then being under the authority would not be under your control. So we should not accept that you consent by receiving benefits. You cannot consent to something without realizing it. So, even though consent might be able to establish authority, some other basis of authority would be required to explain a state's authority over most or all of its citizens.

If we look more closely at consent theory, it normally includes some conditions under which apparent acts of consent are null—without the effect of proper consent. This is a familiar idea: if I threaten to kill you unless you agree to let me use your car, you might say, "OK, take the car." But under that kind of duress the consent has no moral effect. It doesn't permit me to take the car. We could even say that it is not consent at all. In any case it is morally null. But notice what consent theory says about non-consent. Non-consent has a moral force too, of course. For example, if I don't consent to your touching me, then you are not permitted to touch me. But while according to consent theory there are conditions that can disqualify consent, there are apparently no conditions that disqualify non-consent. Non-consent is never null. If there were some conditions that nullified non-consent, the result would be morally equivalent to consent.

This looks worth exploring. What conditions might plausibly be thought to nullify non-consent? Perhaps if the non-consent is morally wrong it should be without moral effect. This isn't plausible in some

contexts. If I refuse to let you use my car, then even if I am wrong to refuse, you still may not use my car. But perhaps things are different with authority, the moral power to require action. Authority is not a form of interference with a person, their body, or their property. It is simply the ability to put them under obligations. So, suppose I propose for you to agree to work under my direction to help clean my garage—a proposal to put you under my authority in this context. And suppose you would be wrong not to agree to it, possibly because I have given you lots of help lately, in addition to really needing your help now. Should we think that you can escape the obligations of authority by violating this obligation to agree to the authority? We might say, instead, that you are under my authority because you would be morally wrong to refuse to consent—call this normative consent. It is hypothetical: you would have consented if you acted morally correctly when offered the chance to consent. This would simply mean that you would be obligated to do what I asked of you. It would not mean that I could coerce you to do it, but only that you would be obligated to help, under my authority, even though you haven't consented to that authority. I find it difficult to see what moral reason there would be for thinking that you could escape the obligations to obey me by immorally refusing to consent to my authority. I will suppose, then, that normative consent is a basis for authority.

THE JURY ANALOGY AND THE COMMITMENT TASK

Using the framework of normative consent in order to explain democratic authority (all within the structure of epistemic proceduralism) gives us the following question: would people be morally obligated to consent to the authority of an epistemic proceduralist democratic political system? If so, why? We noticed that a jury trial nicely illustrates the epistemic proceduralist structure of authority. I want to appeal to the jury trial again now, for a slightly different purpose. If we can see why people would be obligated to consent to the authority of a jury system, I believe that a strikingly analogous argument suggests itself for thinking that people would be obligated to consent to epistemic proceduralist democracy.

Authority is the moral power to require or forbid action. How does a system of jury trials have that kind of authority? When there is an adequate public system of justice, private punishment is forbidden,

whereas in the absence of such a system at least some punishments for some wrongs would be permitted. Vigilante justice is commonly assumed to be morally wrong once there is an adequate public justice system. Why would it be wrong for someone, if asked, to refuse to consent to a jury system's having authority?

The absence of public criminal justice, the world over, would be a great humanitarian problem, a problem on the scale of, say, world poverty. We all have obligations to help solve these problems. It isn't clear how much this demands of us, but few will doubt that we are required to do something. It is plausible to think that the best, or at least a good, solution to the problem of public criminal justice is one in which people are covered by local justice systems—in districts, so to speak. Each is obligated by the commands of the local district, say one's state, but the obligation stems from this being a sensible solution to the larger global humanitarian need for public justice.

Even if we were persuaded that there were obligations to promote a districted public justice system of this kind, that wouldn't yet yield an obligation to do as you're told by your district. I want to grant that authority of this kind raises a special burden for justification: there must be some link to the obligated person's will.

The key is to distinguish between two pressing tasks: the one we have mentioned, which is the need for an adequate and effective public system of justice; the other is one we haven't yet mentioned, what I take to be the pressing task of having people committed to obeying the public justice system. Their being committed to obeying isn't the same as their obeying, but it is a separate important task, something highly valuable in its own right. We often want something morally committed to us even apart from whether that will help us obtain it. Wedding vows might be like this, or promises to repay debts to friends. The commitment task, the task of having people obligated to obey their district for the administration of public justice, is an important task in its own right, important enough that each of us would be wrong to refuse to commit if offered the chance. Consenting to the authority of that system would be required, and so there is normative consent. Normative consent, then, establishes the system's authority. After that, the duty to obey is not directly based on its being a contribution to any important task, but stems simply from the obligation to obey the criminal justice system's verdicts.

The parallels between the jury case and the larger democratic case are very strong. The essential elements of the argument for the authority of

the jury system are all present in a democratic system of government. First, there is a very great value, one that no qualified point of view could deny, to having laws and policies that are substantively just. Second, a proper democratic procedure, like a jury, is (or can be) demographically neutral, blocking the qualified objections that would be possible to any invidious comparisons as between the supposedly wiser and less wise citizens. Third, a democratic procedure involves many citizens thinking together, potentially reaping the epistemic benefits this can bring, and promoting substantively just decisions better than a random procedure. So, fourth, I conjecture there is no nondemocratic arrangement that all qualified points of view could agree would serve substantive justice better. In light of all this, citizens would be morally required to consent to the new authority of such a democratic arrangement if they were offered that choice. Non-consent would be null, and so the fact that no such consent is normally asked or given makes no moral difference, and so any existing democratic arrangement that meets these conditions has authority over each citizen just as if they had established its authority by actual consent.

UTOPOPHOBIA

Thinkers about politics are, for some reason, more concerned with “realism” than are thinkers in moral philosophy generally. In an effort to avoid “utopianism,” it is very common to see fundamental normative standards adjusted so that there is some reasonable likelihood that they will be met in practice, and no similar tendency to dumb down moral principles. Moral philosophers know that people are likely to lie more than they morally should, but this doesn’t move many theorists to revise their views about when lying is wrong. Things are often different in political philosophy. So, for example, many democratic theorists think that standards of political legitimacy should not depend on citizens doing much more than looking out for their own interests in a pretty casual way, and they often think this precisely because they think that is how people are likely to act.

Epistemic proceduralism, of course, suggests that the casual pursuit of self-interest would not be enough. Just as the jury system would lack the epistemic value on which its authority depends if jurors devoted little effort to reaching good verdicts, no democratic system made up of predominantly selfish uninformed voters would have the epistemic

features I have been saying their authority is based on. Does this mean epistemic proceduralism is objectionably unrealistic? If the charge is that practice is not likely to live up to the asserted moral standard, then there are three natural ways to reply.

First, there is the familiar charge that voters are and always will be woefully ignorant and selfish. Great portions of the electorate are ignorant of basic facts about the political system, who holds important offices, which candidates would favor the same things they favor, and so on. It helps in putting this kind of data in context to know that parents, when polled about important matters pertaining to raising healthy and educated children, perform pretty poorly. There are good questions about how they could make good decisions without being able to do well on questionnaires, but this is hardly an absurd possibility. What about voters? There's no reason to be complacent about the state of voter competence, but we should be reluctant to infer from voters failing these quizzes to the conclusion that they are incapable of making good decisions.

Consider, next, the possibility that the moral standards should be weakened to accord better with what can be reasonably expected in practice. I treat this together with the third gambit, the suggestion that it is no flaw in a theory to have standards that are unlikely to be met in practice. To weaken what we take to be the appropriate moral standard we would need a moral argument. "That standard is not likely to be met," does not invoke any moral consideration at all against the standard's truth. Why should it make us think that legitimacy or authority require less than we had first thought?

Suppose the challenge were sharper. What about, "It is not a standard people are capable of meeting"? There is a very tempting but unwarranted slide that often happens from "You and I both know that will never happen," to "That's impossible." It is plausible that if people are incapable of doing something then they are not responsible for not doing it, but the mere improbability of your doing something does not insulate you from responsibility in the way your being unable to do it would. Some things that we all know you will never do are, nevertheless, not impossible, and not even the least bit difficult. It is pretty easy to dance like a chicken in front of your boss. Put your hands up under your arms, thrust your head forward rhythmically, and so on. It is easy, but you and I both know you will almost certainly never do it. The same goes for certain things that might be morally required. Maybe you and I both know that you will not tell your mother that you love her. But that

does not show that you can't do it, or that it is so difficult that you are not responsible for it if you fail. You could certainly do it, you just are not likely to. That fact, that you are unlikely to, is not even the beginning of an excuse.

I want to pause to say, loud and clear, that I am not conceding that what is needed by epistemic proceduralism is highly unlikely, much less certain never to happen. I am just unsure about that. Rather, I want to ask, what if that were so? Would it be a devastating objection to epistemic proceduralism? It would not. If utopianism is the defense of political standards that are very unlikely ever to be met, it is hard to see why it would be a vice, or why political theorists should be so in the grip of what we might call *utopophobia*—the fear of normative standards for politics that are unlikely ever to be met. (There's no similar epithet in moral theory generally, is there?) Normative standards that people are incapable of meeting are much more dubious, so what the critics of supposedly "unrealistic" normative theories need to show is not that "you and I both know it will never happen." That's no objection to a moral theory of politics. They would need to show that not only will it never happen, it is not something people *could* do (or, at least, not without more effort or sacrifice than it's appropriate to require). Maybe epistemic proceduralism asks more of voters than they will ever deliver, maybe not. Either way, this is no deficiency in the theory whatsoever.

We will treat the charge of voter selfishness only very briefly. The first point is that empirical studies of the question have had a hard time verifying the widespread view that voters are largely moved by their own perceived interests rather than by more agent-neutral values.⁴ A little folk experiment is suggestive here. I often ask my students two questions. First, which is it: do most people vote selfishly, or more for the common good? Overwhelmingly they say people vote selfishly. Then I ask, what about you? Overwhelmingly, they say they vote for the common good. Are they being dishonest? They have little to fear from being honest. I do not take their names when they raise their hands to answer. Are they deluding themselves? I see no reason to think so. Is this an artifact of the narrow sample—college students? I doubt that, but I invite you to try it on other groups. Are they biased in favor of themselves without knowing it? Well, I think they probably do know it. We all know that, and so self-interest and other less-than-impartial concerns play an important role. But this simple experiment should shake up any easy confidence that voters are overwhelmingly out for themselves. Hardly any voters see their own motives in that way.

Realism is a vague and dubious constraint when the question is a moral one, when the question at hand is what is right, or just, or legitimate. Obviously, we want to avoid falsehoods. But this includes falsehoods about what bearing people's likely behavior has on what moral standards apply to them.

POOLING WISDOM?

If you have 1,000 coins, with each one slightly weighted to turn up heads—say with a 51 percent chance—what is the chance that at least a majority of them will turn up heads? With that many coins, we know that very nearly 51 percent of them will turn up heads, and so it is quite likely indeed that more than 50 percent will. So now, suppose that, rather than coins, it is 1,000 people, facing a true/false question. And suppose that each person has about a 51 percent chance of getting the right answer (suppose, if you like, that the question is drawn from a pile of which each knows exactly 51 percent of the answers). What is the chance that at least a majority of them will get the right answer? Again, the likelihood is very high, because it is almost certain that about 51 percent will get the right answer, and even more certain that at least 50 percent will. So, under those conditions, the group under majority rule is almost certain to get the right answer. The mathematical fact behind this fascinating scenario was first proven by Condorcet in 1785; it is known as the *jury theorem*, since he was using examples about the likelihood of juries getting the right answer in criminal trials. It has undeniable interest for democratic theory.

If voters are only a little better than random, and choices are between two alternatives, then majority rule would be nearly infallible. Is this the epistemic engine that a theory like epistemic proceduralism needs? I'm afraid that it is not. Consider just a few points.

First, political choices are not always binary, but often take place between several or many alternatives. There will still be some important binary choices: this candidate or that, to build the school or not, and so on. But even if the best choice is likely to be made in those cases, there might be no reason to think that the final two alternatives are the best among the many that were really available. There are some interesting extensions of the jury theorem to more than two alternatives, but the results are not as striking as they are in the binary case.

Second, the jury theorem gives majority rule a high score for accuracy only if individuals are better than random to some significant

degree, not just barely. Our example used 51 percent accuracy for 1,000 individuals, and it wouldn't have worked with only 50.00000001 percent individual accuracy. It's true that the margin above 50 percent that is needed for very high group competence is less if the number of voters is higher, but still, it isn't easy to say what level above random we are entitled to assume. Indeed, and this is the next point, I doubt that we can simply assume that they are better than random at all.

So, third, if you were to ask, "How could a person be dumber than a coin flip?" the answer would be "easily." People have more or less systematic views about many issues. If their system is bad, so to speak, then they could easily be wrong all the time. If, for example, people in some time and place were systematically racist, or sexist, or both, it would not be surprising if their political decisions were worse than the performance of a coin flip would be on political matters involving race or sex. Who knows what other important biases or errors people might have in their systematic thinking on issues?

For these and other reasons, the jury theorem looks like it will not support the kinds of epistemic claims that epistemic proceduralism requires for political legitimacy and authority. There's a second approach, very different but also influential, which we might call the democracy/contractualism analogy. Contractualism, briefly, is a family of views that understand justice or rightness as constituted by facts about what would be agreed to in a certain imaginary collective choice situation. What it is for a law to be just (to give a simplified example) is for it to be permitted by principles that could not be objected to by anyone in an imaginary choice situation in which all participants promoted their own personal (not necessarily selfish) reasons, but no one presses their interests at the unreasonable cost of others. It is not that these imaginary participants take a view about justice, but they do have a limited moral reasonableness to their motives. Justice (or rightness, or whatever—it varies with the theory) is whatever no such person would object to in such an imaginary situation of choice. The results constitute what is right.

If this infallible (because constitutive) procedure is simply imaginary, how is this approach of any use to democratic theory? The reason is that it has seemed to many that actual democratic procedures could look enough like the hypothetical contractual procedure that it might tend to produce the same decisions. The results would not be guaranteed to be just, but a real procedure that resembled the hypothetical procedure might have considerable epistemic value.

The problem with this analogy stems from two facts, one about democracy and one about contractualism. The important fact about democracy is that a unanimity rule is not only impracticable but probably morally inappropriate in any case. Letting a single individual veto proposals that are up for democratic decisions would give too much scope to unreasonable obstructionism, favoring the status quo regardless of whether it was more just than the proposed change. This fact about democracy marks a crucial difference from the hypothetical contractualist situation, whose moral center depends on the veto power. In that morally fundamental imaginary situation, a single victim of a cruel proposal has the power to block it, and if he did not, the contractual approach would not generate plausible moral implications.

That first disanalogy, based on the veto power, might be thought not to be so serious if participants are reasonable in a sufficiently full-blooded way. If each voter, for example, looks out not only for herself but also for others, then the single victim of a cruel proposal might be able to muster enough support to block the proposal even without having the veto power. This will not solve the problem, however, because of a crucial fact about contractualism. Since it is an analysis or explication of the foundations of morality, it would be circular and useless if the imaginary participants themselves already had motives that incorporate answers to the very moral question the contractual situation is meant to provide the answer to, such as "What is right?" or "What is just?" Contractualist theorists recognize this, and they assume the hypothetical participants are addressing some more partial question, such as their own interests, or at least reasons from their own point of view but not from an impartial point of view. What this means is that if actual democratic procedures are going to be analogous, then voters should address some narrow or partial question as well. But if they do, since there is no veto power in real democratic procedures, there is no systematic guarantee that the single victim of a cruel proposal won't ever, or even often, be left out in the cold.

If actual democratic procedures are to have any tendency to produce just decisions it seems likely that this is because participants will have some significant tendency to aim for justice, and not only for some narrower personal ends. Now, a procedure like that might superficially resemble the hypothetical contractual procedure, but it really has no necessary connection to contractualism at all. Whether justice is understood contractually or in some other way entirely, there is some

reason to think that people cooperatively pursuing it might hope to meet with some success. There may be the rudiments of an epistemic approach to democracy in these points, but it owes nothing to a democracy/contractualism analogy.

HOW DEMOCRACY COULD KNOW

Epistemic proceduralism does not need democratic procedures to be highly accurate. This is an easy point to miss, because a natural alternative epistemic approach might say that laws are legitimate and authoritative when they are actually just or correct. Then the general run of laws will only be legitimate and authoritative if the general run of laws is correct—that is, only if the procedure is highly accurate in that sense. Epistemic proceduralism is importantly different. It says that a modestly epistemic procedure gives legitimacy and authority to the general run of laws, even the mistaken ones. The analogy to keep in mind here is the jury trial, since its epistemic value is a crucial reason we think that even erroneous verdicts have authority. That, too, is an epistemic proceduralist structure. The point here is that we are not looking for a source of extremely high accuracy for democratic procedures (though that would be nice). Something quite modest will serve the needs of the theory.

We can start with the very influential idea in recent deliberative democracy theory, that it seems possible to construct an imaginary forum for collective deliberation about political issues in such a way that it would have a strong tendency to make just decisions—to get right answers, so to speak. If there is some feature of actual deliberations that would block this accuracy, we remove it from the imaginary ideal. There are certain familiar features that many think will serve this purpose: all have equal time and power in the deliberation, all address the common good rather than merely some partial interests, all have certain capacities to recognize good arguments against their proposals and others, and so on.

One use of an ideal deliberative situation is to let it actually constitute the truth about rightness or justice. A different use of an imaginary ideal deliberative situation, and the one I propose to rely on, is to treat it as an ideal epistemic situation, not as constituting the truth. That is, this sort of ideal imagines deliberators for whom there are independent facts about what ought to be done. As a consequence, even the ideal

epistemic deliberation can make mistakes. While a morally constitutive ideal deliberation would have to include the veto power, vitiating any serious analogy with democratic arrangements, an epistemic model deliberation has no such need.

This improves the prospects for an analogy between the ideal and actual deliberations, but we should still regard any close resemblance as hopeless. For example, actual political deliberations could not possibly give everyone equal, much less unlimited, time, nor could their rational capacities be what they should be in the ideal. Unlike so many democratic theories that employ a hypothetical ideal (constitutive or epistemic) deliberative situation, the aim here would not be to shape actual institutions and practices in order to structurally resemble the ideal arrangement. If close resemblance were possible that would be fine, but if not there is a “problem of second best”: once certain ideal conditions are violated, it no longer makes sense to think that the other ones are still parts of the second-best scenario. If, for example, one side in a political dispute credibly threatens violence in order to coerce a settlement more to its liking (an epistemically distorting move, to put it politely), what would the goal of mirroring the ideal structure tell us to do? First, of course, it would say to remove that element of force. But suppose that is simply not possible. The choice is either for the other side to threaten some countervailing force, or to stand pat and continue deliberating rationally as it gets politically crushed. Which of these is more likely to produce the same results as the force-free epistemically ideal deliberative situation? In many cases the insertion of additional force is more likely to restore the results to what the model deliberation would have arrived at. This kind of countervailing deviation departs only further from any structural resemblance to the ideal, but since the aim is epistemic there is no reason to seek such resemblance for its own sake.

This model of countervailing deviation from the ideal epistemic deliberation promises to give a more plausible account of what is morally appropriate political behavior—as is too often said, it is not a college seminar. Details about what kinds of political practice would be called for by this model are too sensitive to the complexities of specific contexts to say much useful about in a philosophical treatment. This brings us full circle from our opening disclaimer about the limits of a philosophical framework. I have willfully proceeded at a fairly high level of abstraction, since I believe that this is where many of the most important problems in democratic theory reside. Still, is there no concrete vision of politics that emerges from the distinctive features of epistemic

proceduralism? I conclude with some impressions of my own about implications the view might have in practice, emphasizing that these are not propositions for which I have argued.

First, if points of view get their influence on public conclusions by virtue of the wealth they have at their disposal, public reasoning will be seriously distorted unless this irrational element of power can somehow be countervailed in creative political practice.

Second, legal and social protection for the ability to dissent from orthodoxies and majority positions is not (at least not merely) some right owed to the dissenter, but a crucial ingredient in a healthy public life, one in which there is a basis for hope that the public view might discover and remedy its errors over time and move progressively toward sounder views.

Third, equality in political matters is also not some natural right, even if a certain kind of equal regard is. Political equality depends on, and finds its limits in, what sorts of arrangements will allow the promotion of justice and common good in a way that can be justified to the broad range of points of view that are owed acceptable justifications for the coercive political arrangements under which they live. Inequality of various kinds is bound to pass this test, but I have argued that the overall system seems bound to be recognizably democratic in its procedures for making law and policy. Finally, where epistemic proceduralism's aspirations are met—which might be unlikely, but is hardly impossible—there is an obligation to obey the law. Not just any law, since some could be too unjust or unjust in the wrong way, but including many laws that are indeed unjust. Legitimate politics involves authority, and there is no getting around it. We cannot collectively live as we ought to live and still be under only our own authority. Perhaps this is still all fairly abstract, and, even as far as it goes it is hardly an original vision of politics in modern times. But originality and detail at the level of institutions and concrete practices are not the areas where I have tried to make my contribution. What I offer instead is a philosophical framework.