Michael Bennett

EXPERIMENTS IN DISTRIBUTIVE JUSTICE AND THEIR LIMITS

ABSTRACT: Mark Pennington argues political systems should be decentralized in order to facilitate experimental learning about distributive justice. Pointing out the problems with Pennington's Hayekian formulation, I reframe his argument as an extension of the Millian idea of "experiments in living." However, the experimental case for decentralization is limited in several ways. Even if decentralization improves our knowledge about justice, it impedes the actual implementation of all conceptions of justice other than libertarianism. I conclude by arguing for the compatibility of egalitarian redistribution with the epistemic virtues of markets pointed out by Hayek.

Keywords: *decentralization; epistemic democracy; experiments in living; distributive justice; liberal egalitarianism; libertarianism; Mark Pennington; moral epistemology; rational ignorance.*

It is always a disconcerting experience to read a book that proceeds intelligently, carefully, and in good faith to argue for a political position significantly opposed to one's own. At the very least, the experience should give us pause. This is particularly fitting in the case of Mark Pennington's *Robust Political Economy* (2011), since one of the book's main themes is the importance of intellectual humility in politics. Pennington points out that institutions must be both designed by largely ignorant people, and be designed to cope with people's ignorance. This constitutes one of two main themes of the book, the "knowledge problem." The other main theme of the book is the "incentives problem": institutions have to be able to cope with limited altruism. These two themes represent, to my mind, the most plausible arguments for the political right. As such, Pennington's book should be recommended reading for all open-minded left-wingers.

Nowhere are (at least self-perceived) open-minded left-wingers more concentrated than in the field of post-Rawlsian political theory. In this respect it is worryingly convenient that the "ideal theory" methodologies of political theory make the knowledge and incentives problems invisible. Ideally, everyone would know what they need to know and be motivated to act morally. Implicitly, Pennington seems to grant that under such circumstances widespread government intervention would be desirable—although I am sure he would reject the very idea of abstracting from human ignorance and self-interest as a ludicrous and pointless endeavour. The most persuasive arguments for free markets and minimal states are based on the real world's departures from perfect wisdom and perfect altruism. If

political theorists of the left want to be relevant to actual political disagreements, they must engage on this terrain of non-ideal theory.

I will engage with a crucial element of Pennington's non-ideal case for classical liberalism: his rejection of egalitarianism (and theories of distributive justice more generally). This is crucial to Pennington because, out of the four strongest challenges to classical liberalism, it is egalitarianism that "continues to exert perhaps the greatest influence on public policy debate" (111). Pennington offers several arguments for rejecting egalitarianism, but I will focus in detail on his most plausible argument, which he describes as a subset of the knowledge problem: namely, that egalitarianism is a form of "constructivist rationalism" (119). His argument is that we should design our institutions to improve our understanding of justice. He aruges the institutions which do this best are classical liberal ones that enable people to experiment with a variety of distributive principles.

The importance of this general idea of experiments in distributive justice goes beyond its context in Pennington's rejection of egalitarianism. Pennington is right to think that political theorists have spent too much time advancing their own conceptions of justice, and not enough time thinking about how institutions shape these conceptions themselves. Other theorists have made epistemic arguments in favour of decentralization and limited government, but none that I know of have stressed the element of moral learning in the same way (Brennan 2014; DeCanio 2014; Pincione and Tesón 2006; Somin 2013). If it could be shown that decentralized institutions were better at promoting knowledge of justice, this would constitute an argument for classical liberalism that did not depend on controversial moral premises about natural rights. However, as I will show, we should be skeptical about these claims. The experimental argument for decentralization is limited, and does not point to classical liberalism in the way Pennington imagines.

I start by reconstructing Pennington's argument and criticizing the specifically Hayekian formulation he gives of it. The second section proposes that the Millian idea of "experiments in living" provides a better way of framing the experimental argument for decentralization, although one that is much more limited in its institutional implications. Having set out the potential epistemic benefits of decentralization, the following section looks at the moral drawbacks a more decentralized system has in being less able to actually implement many conceptions of justice. The fourth section criticizes Pennington's equation of egalitarian redistribution with central planning. Redistribution is quite compatible with decentralization justified by the epistemic virtues of markets (as argued for by Friedrich Hayek and others).

The Hayekian Argument for Decentralized Distributive Justice

Pennington's defense of classical liberalism against egalitarianism does not involve a head-on critique of egalitarianism as a moral doctrine.

Instead, Pennington seeks to rely on the "knowledge problem" to show that a state that attempts to pursue distributive justice will be less effective than one that does not. Pennington's argument has two premises. First, given our lack of knowledge about distributive justice, we should design institutions to improve this knowledge. Second, institutions that decentralize responsibility for distributive justice are better at improving our moral knowledge than institutions that allow distributive justice to be pursued centrally by government. I will begin with some brief comments on Pennington's first premise, with which I am largely sympathetic. I will then point out the weaknesses of the Hayekian argument Pennington gives for the second premise.

According to Pennington, taking seriously the limitations of our moral knowledge implies that we:

> must consider what sort of institutional arrangements facilitate the evolution and communication of the appropriate moral principles. Attention should, in other words, shift from a preoccupation with "static" or "end-state" theories to ones which focus on the dynamic processes within which principles of justice may be discovered and disseminated. (121)

So long as we are uncertain about justice, we cannot directly proceed from advancing a set of moral criteria by which society should be judged to recommending the set of institutions or policies that will be optimal according to those criteria. Rather, we should plan our institutions so that our knowledge about the criteria for judging those institutions itself tends to improve over time.

On this important point I fully agree with Pennington. He uses this idea to criticize the kind of political theory practiced by John Rawls and Ronald Dworkin, which proceeds directly from moral philosophy to institutional design. This critique of the philosopher-legislator is reminiscent of Jeremy Waldron's (1999) complaint that contemporary political theorists have spent too much time setting out their own conceptions of justice and not enough time thinking about the processes by which society should choose between competing conceptions of justice. However, whereas Waldron is more concerned with fair ways of resolving disagreement, Pennington's response is closer to that of epistemic democrats such as David Estlund (2009), who are more concerned with how institutions can be designed to better track the truth (whatever that might be).

Pennington argues that the institutions that will best promote learning about distributive justice are those that forbid states from actually pursuing distributive justice in practice. He advocates

> an "open society" that allows different individuals and groups to enter and exit from a variety of cross-cutting arrangements provided by families, firms, neighbourhood associations and the panoply of cultural and civil institutions that distribute money, goods and social status. The processes of interpersonal and intercultural learning and adjustment that would occur in such a framework, though offering no guarantees that the "right" principles might be discovered, would nonethe-

less increase the chance that any person or group picked out at random would have the opportunity to find the "right" set of distributional rules. (127)

The argument Pennington gives for thinking a radically decentralized system will improve learning about distributive justice is a Hayekian one. He seeks to establish an analogy between Hayek's argument for the market and his own argument for decentralized distributive justice. This analogy has two parts: first, that knowledge of justice is relevantly similar to economic knowledge in terms of complexity, dispersion and particularity; second, that knowledge about distributive justice can be communicated by price signals in the same way as economic knowledge. Unfortunately for Pennington, both these aspects of the analogy fail, as I will explain

Pennington seeks to establish an analogy between knowledge of justice and the kinds of knowledge that motivated Hayek's argument for the market. For Hayek, the production and allocation of all the goods in the economy is simply so complex no individual could possibly comprehend all the information at once. Instead, knowledge about production plans and consumption preferences are dispersed throughout the population. Moreover, much economic knowledge takes the form of tacit know-how that cannot be verbally communicated to others. The best economic institutions will accordingly be those that allow decisions to be made in a dispersed way by those who know best, rather than attempting to make all decisions centrally. Pennington tries to apply the same logic to distributive justice: "the knowledge pertinent to appropriate principles of distribution cannot be centralized in a public forum and the complexity of the factors that contribute to the formulation of distributional rules cannot be grasped by a group of minds engaged in a deliberative exercise" (126).

Pennington's analogy fails here because knowledge about justice and economic knowledge are relevantly different. Distributive principles like "maximise aggregate welfare" or "maximise the primary goods of the least advantaged" are not necessarily *complex* at all. Economic knowledge involves (often tacit) knowledge of particulars, whereas knowledge of distributive justice (at least in the work of philosophers like Rawls, Dworkin and Robert Nozick) is the knowledge of general rules. In his classic article on the subject, Hayek suggested that the element of particularity in economic knowledge was crucial to the epistemic case for markets: "If it is today so widely assumed that the latter [central planning] will be in a better position [than the market], this is because one kind of knowledge, namely, scientific knowledge, occupies now so prominent a place in public imagination that we tend to forget that it is not the only kind that is relevant" (Hayek [1945] 1949: 80). The implication seems to be that general rules are not subject to the epistemic case for the market.

If distributive principles are simple and general in the way envisioned by Rawls, Dworkin, Nozick and others, they can be "centralized" in the mind of a single individual without loss. They are thus immune to the Hayekian critique of "central planning" that Pennington seeks to apply to them. This should not be taken to imply a more

general rejection of the epistemic argument for the market in the context in that Hayek originally made it. Even if distributive principles themselves are simple and general, knowledge of *how* to best follow distributive rules is complex, dispersed and tacit precisely because it also requires economic knowledge. For this reason, the best way to fulfill a principle like utilitarianism or the difference principle may well involve markets. Hayek's argument for the market can be read as a form of indirect utilitarianism along these lines. I am not disputing epistemic arguments for markets in general (I will return to this subject below in the fourth section). My point is that it does not work in the specific case of knowledge about distributive principles.

The second element in Pennington's analogy with Hayek concerns the role of competition and inequality in signalling. For Hayek, profits and losses inform entrepreneurs which business models are most socially useful and so should either be emulated or avoided. Pennington thinks a similar process of signalling will take place with distributive norms:

> *It is differences in results - that is, inequalities - that enable decentralized agents to learn about the merits and demerits of alternatives, and for these results to be disseminated across the overlapping perspectives of countless actors none of whom may be aware, individually or collectively, how their particular choices contribute to the evolution of distributional norms (126).*

This analogy between product markets and competing distributive norms is puzzling. There is no reason to think that material inequalities between communities will be a good signal of the merits and demerits of different distributive norms, unless we simply assume that prosperity is the correct moral criterion. Pennington rejects the idea (which he attributes to Dworkin) that all relevant goods should be denominated in money prices, and stresses the importance of other forms of competition, such as competitive federalism. However, he maintains that competition and unequal results provide useful signals to others about the normative validity of distributive principles, and this is problematic.

Elsewhere in the book, Pennington explains his evolutionary approach to economics as an alternative to neoclassical general equilibrium analysis. He considers the worry that the survival of the fittest is normatively uninteresting unless there is some reason to think that "fitness" is normatively valuable in a particular context. In the case of the market, the "selection criterion" of fitness is the "account of profit and loss," which is determined by firms' responsiveness to consumer choices (43). The discipline of the market is thus normatively defensible ultimately because it tends towards the efficient satisfaction of consumer preferences. While this response may be satisfactory as a defense of commodity markets, it will not do at all as a defense of radically decentralizing distributive justice. Consumer sovereignty is not the right kind of determinant for the success of moral principles. The point of a principle of justice is not to efficiently satisfy consumer preferences, but to do the morally correct thing.

Estlund (2009, 177) helpfully presents Hayek's epistemic approach to the market as a case in which the answer to one question (what is the efficient allocation of resources?) is provided by entrepreneurs separately answering a different one (what will maximize my own profits?). However, as Estlund points out, this sort of indirect question-answering is surely the exception rather than the rule. Employing this kind of indirect method to get at the truth requires a special justification. Hayek's epistemic argument for the market can be seen as such a justification, because he stresses the role of tacit, dispersed knowledge and because he implicitly assumes that the criterion for economic success is the efficient satisfaction of consumer preferences. However, when the criterion for success is itself the thing we want to discover, there is no reason to think this particular justification for markets applies.

Pennington's Hayekian argument that decentralization promotes knowledge of justice fails. However, the goal of promoting knowledge by means of institutional design remains a valid one. And even if Hayekian competition does not work, there might be other mechanisms by which decentralization can promote moral knowledge. In the next section, I will explore these possibilities.

A Millian Argument for Decentralized Distributive Justice

The most promising source from the liberal tradition for thinking about decentralization and moral learning is not Hayek but John Stuart Mill. In this section I will consider whether the Millian idea of "experiments in living" provides an epistemic reason to decentralize distributive justice. Rather than trying to be faithful to Mill's text, I will follow Ryan Muldoon (2015) in seeing how far the idea of experiments in living can be extended for use in contemporary political philosophy. I will start by outlining a Millian argument for decentralization before considering two problems with the argument: an objection from moral epistemology and an "incentives problem" of the kind familiar from other parts of Pennington's book. In the following section, I will consider a deeper problem with the approach, which is that it requires some people to accept unjust institutions in the medium term as part of the experiment.

The idea behind experiments in living is that trying out different modes of life improves our knowledge about how to live. Applied specifically to distributive justice, the idea would be that experimenting with different distributive principles would improve our knowledge about what justice requires. It is important to distinguish the concept of experimentation used here from the concept of experimentation Hayek applies to entrepreneurship. In the Hayekian framework I considered in the previous section, the idea of experimentation is bound up with the idea of competition. Entrepreneurs are imagined as experimenting with different products, with the account of profit and loss providing the results. In the ordinary scientific usage of "experiment," however, there is no necessary connection with competition, and the same is true of Millian experiments in living.

Individuals and communities experimenting with different modes of life are no more in competition with one another than sets of siblings are in competition with one another over the best ways to care for elderly parents. A successful experiment by one party is in no way unwelcome to other parties. Experimenters interact through emulation, but the pressure to emulate success is nothing but the internal desire to do the best oneself. It is quite different from the external necessity of keeping up with best practice to avoid bankruptcy, which is enforced by competitive markets. A successful experiment is a pure positive externality for those who learn by observing it.

If the link between experimentation and decentralization has nothing to do with competition, why would we expect a more decentralized political system to better facilitate experiments? Pragmatist political theorists argue that successive governments in a unitary political system can be viewed as different "experiments" from which we learn each time (Anderson 2006; Gerlsbeck 2016; Knight and Johnson 2011). However, the advantage of decentralization (seemingly missed by these pragmatists but pointed out by Pennington) lies in being able to run several different experiments simultaneously. This creates many more data-points than would be possible by running experiments sequentially in a centralized system, and permits them to be compared more systematically. Running more experiments creates a bigger and thus more reliable body of evidence from which to draw conclusions.

What are the institutional implications of this argument for decentralization? Because Millian experiments have no connection with markets, we should not assume that "decentralizing" justice on this basis will take the radical form of marketization that it has in Pennington's Hayekian argument. Muldoon suggests the Millian argument implies some presumption in favor of individual liberty against the authority of the democratic state. However, this is not the only institutional way of decentralizing the pursuit of justice. An alternative approach is Keynes's (1933) suggestion that different countries should experiment with different institutions. Differences between countries are welcome, we might say, because those differences allow us to learn something about the different values that those countries' institutions are oriented around. Pennington expresses support for the idea that countries should experiment with different development strategies (216). However, the Keynesian approach is sharply distinguished from the inter-jurisdictional *competition* that Pennington also favors (59). Unlike efficiency-based arguments for competitive federalism (Tiebout 1956), the experimental argument for political decentralization is not necessarily tied up with the unrestricted freedom to enter and exit different communities. Free movement of people might be epistemically beneficial if voluntary participants in experiments are more enthusiastic and so provide a better guide to the desirability of a society guided by a particular principle of justice. But if we want to be genuinely impartial about which principles of justice are worth exploring, the logic of the argument leads to

the conclusion we cannot rule out experimenting with illiberal principles of justice which would deny individuals the right to exit.

I will now consider a potential problem with the moral epistemology that is implicit in the Millian argument. The benefits of experimentation are relatively obvious for questions of instrumental reason. For example, being able to observe countries governed by a variety of different economic institutions is very helpful for deciding which institutions will do best by a criterion such as aggregate welfare or distributive equality. The idea of using real-world experiments to settle non-instrumental questions is a little more contentious. Mill's own text is unclear as to whether he thought experiments in living would shed light on the best way to realize abstract goods of human flourishing that we already think are valuable, or whether experiments could also help us discover which abstract goods are valuable in the first place. (Keynes is similarly unclear). In contemporary philosophical parlance, the question is whether experiments improve knowledge only of conceptions of the good, or also of conceptions of the right (justice). It is the latter kind of knowledge that is necessary for the argument for decentralization considered here.

The epistemology of learning about the nature of justice through real-world experiments is controversial. One way of defending the Millian argument would be to simply reject the common idea that normative principles must be ultimately independent of empirical facts (Cohen 2008, ch. 6; Miller 2008). However, such a move might not be necessary. Knowing certain facts about the world may be epistemically useful for discovering moral principles, without it being that case that the truth of the principles logically depends on those facts. Marc Stears (2005, 333–35) raises this point in his discussion of G. A. Cohen's fact-insensitivity thesis. That normative justification is ultimately grounded on fact-insensitive principles does not mean that our knowledge of those principles is similarly independent of empirical facts. Muldoon makes a similar point in raising the distinction between "discovery" and "justification": "How one comes up with a new idea is entirely separate from whether that idea can be justified" (10).

By what mechanism might this work in practice? One possibility is that experiments sharpen our intuitions. If our judgments are more likely to be correct when we consider actual moral dilemmas rather than mere thought experiments, this would be a reason to want more actual experiments in the world. Muldoon seems to pointing to a different possibility, however: that experiments will be useful in prompting people to ask new moral questions and adduce new kinds of reasons as moral "evidence." Some mechanism of this kind will have to operate if experiments in distributive justice are to be informative.

Assuming that some mechanism for moral learning by experiment exists, it will surely be much less efficacious than experiments in science or business, for two reasons. First, experimentation in justice will necessarily be a much slower process than consumers experimenting with different products in markets. Consumers can try out a range of different products at once, and they can even try them out side by side. A reliable experiment on a principle of justice, on the other

hand, requires a body of participants who will sincerely commit themselves to the relevant way of life over the long run. Only then will we get an idea of what a community run by such a principle would look like. Second, observing and emulating other ways of life is much more difficult than producers observing and emulating alternative production techniques. Only participants themselves get a full idea of what living under a society ordered by a particular principle of justice feels like, because truly understanding a norm requires actually following it and making it a constitutive part of one's own identity. Observing communities structured by alternative moral norms may offer some provocation to change our moral beliefs. More often, however, we will simply apply our existing moral beliefs when we look at other communities. Thus, even if the moral-epistemology problem with experiments in distributive justice can be surmounted, the case for decentralization that emerges is weak.

The second major problem with the experimental argument for decentralization concerns its robustness in circumstances of limited altruism. If we grant that political decentralization makes institutions more robust to the knowledge problem, Pennington's avowed program also entails analyzing how robust these institutions are to the incentives problem. Millian experiments in distributive justice are not at all robust in this respect, because personal prudential incentives run directly counter to the interests of the system. Experiments in living are robust to limited altruism when it comes to conceptions of the good (the choice of a religious or an artistic life, for example) because if experimenters discover new and valuable ways of life, most of the prudential benefits of this are internalized by individuals themselves. However, , if experiments improve our knowledge of justice, the prudential benefits of this are almost entirely external to the experimenters. Faced with a choice of communities following different distributive principles, many people will choose not the community whose rules they judge to be most just, but the community whose rules favor their own self-interest. Thus, Wilt Chamberlain is likely to be strongly tempted by a Nozickian community, whereas the person with congenital health defects is likely to be tempted by a Dworkinian community. Moreover, as soon as some people become start behaving in this way, others are likely to follow suit, because willingness to incur sacrifices for the sake of justice is often conditional on a perception that other people are doing the same. Whereas the Hayekian version of experimentation works along the grain of prudential incentives, Millian experiments in distributive justice work against it, requiring people to be independently motivated to search for the truth about justice even at cost to themselves.

One way of dealing with these incentive problems would be to restrict freedom of movement between communities experimenting with different norms. At the extreme, simply forbidding movement would remove the possibility of choosing distributive communities for self-interested rather than moral reasons. Setting this aside as unacceptably illiberal, we can pick out large-scale movement of capital as a particular target for restriction. Whereas the movement of individuals

and their modest holdings are plausibly driven by the search for a morally superior living environment, movements of large fortunes tend to be driven by self-interest. As Keynes pointed out, the free movement of capital tends to be a force for international homogeneity, limiting the scope for diverse distributive experiments.

Pennington agrees that, on the consumption side, free trade reduces international diversity while increasing intranational diversity. However, he also claims that free trade has not promoted cultural homogeneity on the production side, because it allows countries to develop comparative advantages in sets of norms that are suited to different products (Japanese-style teamwork for car manufacturing, Chinese-style entrepreneurship for fashion) (213–14). This seems doubtful to me, but it merits empirical investigation. Pending this, we can say that to the extent that free movement of capital is a force for homogeneity, the experimentation argument provides reasons for measures such as capital controls and withholding taxes, which minimize the extent to which investment flows are driven by differences in regulatory environments (Avi-Yonah 2000). Once we tackle the incentives problems with experiments in distributive justice, the kind of decentralization that this argument justifies is very different from the one Pennington envisages.

The Moral Costs of Decentralization

The previous section argued that decentralization had some tendency to advance knowledge about justice, but that this advantage is likely to be quite weak. However, something is better than nothing. Unless the decentralization proposal carries other costs, it will still be worth pursuing. This section will focus on one such cost: that, according to many conceptions of justice, decentralization stands in the way of the policies that would promote justice. That is, institutions that improve knowledge can also impede the ability to act on that knowledge. Once we take account of these costs, the moral impact of decentralization may be net negative.

Knowledge and implementation can come into conflict in several ways. King Solomon's proposal to cut the baby in half was a good way of finding out who should have been given custody; but if actually implemented, it would also have rendered that knowledge useless. The economics of search theory illustrate a related dilemma: How much money should you spend trying to acquire valuable empirical private information that you know in principle to exist? Tradeoffs between knowledge and implementation become even trickier when the knowledge involved is moral rather than empirical in nature, or when it involves acquiring new concepts rather than acquiring new data-points of a familiar kind (on this distinction see Fearon 1998, 49–52).

The task of deciding what criteria should be applied to oneself has a reflexive element, because the resources spent on reflection about what criteria to apply might themselves be judged to be too large or too small. However, this reflexivity is not particularly compli-

cated or problematic, and it is a phenomenon we are very familiar with in the case of individuals. Consider someone deciding what their goals should be for the day, in a case where there are better or worse answers to the question. The more time they spend thinking about the choice, the more likely they are to get a good answer. However, if they spend all day asking what they should do, they will have no time left to actually *do* it. So (much of) their time spent thinking will be wasteful, according to the criterion decided on by that thinking itself. It is clear that the two extremes of all thought, no action and all action, no thought are both unlikely to produce good outcomes. However, when knowledge is moral and conceptual in nature, I cannot imagine how the decision to stop thinking and start acting could be anything other than a matter of contextual practical wisdom that cannot be codified.

The tradeoff between knowledge and action that is familiar for individuals also applies to institutions. Principles of justice are the criteria by which institutions are evaluated; but, as Pennington points out, institutions can also affect our knowledge of these principles. One obvious way this tradeoff arises in institutional design is in considering how much time and resources should be devoted to democratic deliberation. Would the epistemic benefits of paying all citizens to skip work for a "deliberation day" before elections (Ackerman and Fishkin 2002) exceed the economic costs? Insofar as the experimental argument for decentralization succeeds, the optimal extent of political decentralization is another case of a tradeoff between knowledge and implementation.

To see this, consider a choice between two institutional systems: a conventional unitary democracy *versus* the decentralized classical liberal system proposed by Pennington, in which the state is restricted to enforcing property rights. Suppose, on the basis of the argument in the previous section, that the classical liberal system is epistemically superior; people in the classical liberal system tend to have more accurate beliefs about justice. Now consider two possibilities. First, it might be the case that the correct conception of justice is, in fact, a libertarian one. In this case, all is well for the classical liberal system; it is better at discovering the truth about justice, and moreover it has, in effect, already implemented it. Second, consider the possibility that the correct conception of justice is some form of distributive egalitarianism. In this case, by assumption, people in the classical liberal system are more likely to come to favor egalitarianism than people in the conventional unitary democracy because they can experimentally see egalitarianism in action. However, implementing the egalitarian principle of distributive justice will be much more difficult in the classical liberal system than in the conventional unitary democracy. The democracy requires only a majority to implement a distributive principle. The classical liberal system, in effect, requires unanimity. Indeed, exactly the same issue will arise for *any* principle of justice other than a libertarian one. Except in the special case of libertarianism, the classical liberal system is much worse at implementing principles of justice, even if it is better at discovering them.

If the classical liberal system is unable to implement non-libertarian society-wide principles (such as egalitarianism), it is also questionable how far it is actually able to experiment with such principles in the first place. In the classical liberal system, it is possible to experiment with egalitarianism and other principles at a smallscale. However, the knowledge gleaned from observing such small-scale experiments may be of only limited relevance for considering the merits of applying the same principles at a much larger scale.

When Pennington asks "what the requirements for a genuinely 'impartial' choice of social institutions might be" (127), he makes no mention of this problem. Yet the classical liberal system he favors is not actually impartial between different conceptions of justice. Classical liberalism appears to offer impartial terms to all distributive views because it enables any group of people to form their own distributive community. However, this sort of impartiality will only really satisfy libertarians (who are happy that the system as a whole respects ownership rights) and cultural relativists (who only wish to apply their principles in their own immediate neighbourhood). Impartiality of this sort does not take seriously theories such as those of Rawls, Dworkin, and Iris Marion Young-theories to which Pennington pays a great deal of attention, but which apply to society as a whole. Rawls, Dworkin, or Young would not be satisfied by the thought that their own communities were just if those communities were merely islands in a larger sea of injustice. These theories are intended to describe the demands justice makes on everyone, not merely the demands it makes on those agree with the theory.

That Pennington regards contrary conceptions of justice as merely private preferences is indicated by his statement that

> even in the case of what are more explicitly moral questions, reliance on the exit principle may be preferable. Those who find themselves in the minority on whichever side of the abortion debate, for example, may find it better to live under a regime of interjurisdictional competition where they have at least some scope to live under an authority that reflects their own point of view, rather than face a unitary system of moral laws. (76)

Pennington seems to assume that moral beliefs are desires to live in particular kinds of communities rather than imperatives to make the world a certain way. To favor interjurisdictional competition and open borders is effectively to privilege the pro-choice side of the debate, since it means that women even in the pro-life jurisdiction can always receive an abortion if they are willing and able to pay the costs of exit. At the same time, interjurisdictional competition does not fully address the concerns of feminists, who do not just want access to abortions themselves, but want women everywhere to have access without having to pay those exit costs. A liberal approach to abortion may well be justified on directly moral grounds. But classical liberalism cannot avoid taking a stance on abortion as a moral issue once we recognize that it is indeed a moral issue rather than merely a matter of personal

preference. The same applies to wider moral issues, such as questions of distributive justice.

A genuinely impartial mechanism for institutional choice must be one that confers no substantive advantages or disadvantages on any particular conceptions of justice (within some suitably wide boundary of reasonableness or tolerability [cf. 128]). Democracy is just such a mechanism, because it can implement any distributive rule, including a libertarian one.

In saying that, I have no desire to reject the entirety of public choice theory and its associated critique of democracy. Public choice theory can help to identify biases in political systems while remaining impartial between different conceptions of justice. As far as possible, political systems should be designed to counter these biases. The problem of rational ignorance, which Pennington emphasizes (61–69), is a good example of this, and is worth briefly dwelling on (see also Caplan 2011 and Somin 2013).

The theory of rational ignorance holds that voters in large electorates, recognizing the meager odds that their votes will (individually) tip the balance in an election, decide to underinform themselves about politics and government, since becoming well informed would almost certainly be a waste of time and energy. Let us assume that this is a genuine problem. It may just be a problem we have to live with if we cannot come up with any impartial alternatives. This point is structurally similar to Pennington's complaint that certain so-called market failures are not really failures at all if no alternative institutions would be able to perform better (22–40). It may be that we face a trade-off between market systems, in which citizens are well-informed but selfish, and voting systems, in which citizens are morally motivated but poorly informed. The low probability of making a difference which makes it rational for voters to be ignorant, also makes it rational for them to vote in a moral rather than a selfish way, since doing so costs them nothing.

Less pessimistically, there might be institutional innovations that avoid this trade-off without abandoning the basic principle of democratic sovereignty. One proposal, advanced by Claudio López-Guerra (2011) and others, is to make greater use of lotteries rather than elections to select decision makers (Guerrero 2014; Landemore 2013, 108–17). A lottery could be used to select members of an upper house of parliament, or to narrow the franchise for parliamentary elections. With a small group of randomly selected voters, it would be financially feasible for them to undergo some sort of educative deliberative process before casting their votes. Reducing the size of the decision-making body would also change the incentives that create the rational-ignorance problem. Individual votes in a smaller body would have a much greater chance of altering the outcome, thus increasing the incentive to acquire and rationally process information about the decision. At the same time, increasing the efficacy of individual votes would be unlikely to raise the alternative problem of selfishness. Because they lack any common social features, lottery-selected decision-makers would only be able to enrich themselves by increasing the

perks of their own specific offices. Such nakedly venal conduct would encounter resistance from the consciences of the decision-makers themselves, the opprobrium of their fellow citizens, and (if necessary) constitutional provisions enforced by another branch of government.

I am not suggesting that democratic innovations like these can remedy all the problems with democracy to which Pennington and others have pointed. I fully agree that citizens should be aware of the knowledge and incentive problems with democracy, and should accordingly exercise restraint and humility in directing the state to intervene in social life. This may well entail that democratic states choose to leave many areas of life up to markets and private initiative. In the next section, I will argue that this sort of humility is nonetheless quite compatible with a strongly redistributive state.

The Possibility of a Liberal Egalitarianism

So far, I have argued that the negative argument Pennington makes against state distributive justice is unsuccessful. In this section, I would like to allay some fears about distributive justice and its compatibility with the epistemic merits for which classical liberals such as Hayek have prized markets. Part of Pennington's rejection of distributive justice seems to be based on viewing redistribution as a form of central planning, but there is no necessary connection between the two.

One argument Pennington repeatedly uses against government intervention is that even if free markets are producing suboptimal results, there is no reason to think government will have the knowledge necessary to adjust markets in the precise way necessary to produce optimal results. This argument is invoked against interventions to remedy information externalities and lock-ins (33, 37), to provide public goods in health and education (167), to deal with environmental externalities (229, 241), and to promote distributive justice (133–34). This reasoning involves the same kind of mistake Pennington describes as the "nirvana fallacy" earlier in the book (24). The nirvana fallacy is to take the inability of markets to achieve perfect Pareto optimality as a reason for states to intervene, without considering whether state interventions might actually make things worse. Similarly, Pennington complains that state intervention fails to achieve optimal results without allowing that suboptimal intervention might nonetheless be better than no intervention at all. Governments do not need to hit the perfect optimum of distributive justice or efficiency; they just need to make things *more* just or efficient than unconstrained markets would. We don't need to know what the ideal level of public spending on roads would be in order to say that some public roads would be better than none public roads at all (cf. Sen 2006). Clearly, beliefs about which direction to adjust in—more roads or or less?— implies beliefs about where the optimum level is *not;* but it is compatible with considerable uncertainty about where precisely the optimum is. Just because concerted public action will not hit the optimum, it does not follow that public action is unable to approximate towards

it—at least compared with leaving things completely up to an uncon-strained market.

Pennington's application of this nirvana fallacy to distributive justice in particular shows an ironic distrust of markets and spontane-ous orders. He writes that

> *to suggest that a particular decision increases or decreases social jus-tice,the relevant actor, whether a private individual (through a chari-table action)or a government department (through tax and spending decisions) would need answers to the following kinds of questions, among many others. Will the recipients of transfer payments engage in activities and spend the relevant money in ways that increase or de-crease their earning power in the labour market? To what extent might transfer recipients spend on a consumption bundle which increases the demand for unskilled labour, and thus reduces income inequality rela-tive to spending, that raises the demand for already higher-paid occu-pations and thus may increase inequality? [The list goes on.] (134)*

If we really believe in the virtues of markets, we don't need to answer all these questions; we simply trust that the market will make the necessary adjustments. If recipients of redistribution choose to use it to work less or to consume particular goods, then if we really believe individuals have a better idea of their own interests than experts do, why should we complain? Pennington's general case for the market, with which I sympathize, is that is *robust*: a decentralized network is much better at adapting to external shocks than a centralized hierar-chy. It thus seems odd that Pennington thinks that markets are so fragile when confronted with state intervention. This is not to deny that there are debates about tax neutrality and about the relative merits of income, wealth, inheritance, and consumption taxes. How-ever, this debate is relatively trivial and academic compared to the broad sweep of Pennington's concerns. Each of these taxes, applied generally, leave investment, production and consumption decisions in the hands of the market rather than the state.

Redistribution is not central planning, and it does not require the kind of homogenous agreement on ends that Pennington finds so troubling. Once we distinguish between the right (how resources and opportunities should be distributed) and the good (what people should do with them), we can demand that society follows certain principles of right while leaving questions of the good up to individuals. Penning-ton's concern for toleration and allowing others to pursue different ends of their own choosing (91) is fully compatible with a liberal egalitarianism. Redistribution need not reduce the net amount of individual freedom in society. Instead, it merely shifts around the freedom that property confers (Cohen 1994, 14–16). Moreover, it can do so in a predictable, rule-bound way that avoids the pathologies of arbitrary rule that so concerned Hayek. Most redistribution in contemporary welfare states takes place through in-kind or conditional benefits. However, if in-kind or conditional benefits are judged to be problematic for knowledge or incentive reasons, redistribution can always take place instead through cash, either upfront as a stakeholder

grant, or at regular intervals as a basic income (Meade 1964; Van Parijs 1997). Redistribution, in other words, can take place in a free-market society.

I wish to conclude by pointing out the limits of my critique, and asking where it leaves Pennington's argument for classical liberalism.

My critique is tempered in two ways. First, I have shown not that there is no knowledge-based reason for political decentralization, but that the case for it is much weaker than Pennington supposes. The epistemic benefits of political decentralization for distributive justice must be balanced against the disadvantages political decentralization has for the implementation of almost all distributive principles. On balance, the case for decentralization is probably quite weak, especially at the intrastate level. In practice, its implication is probably a consideration in favor of countries or federal jurisdictions being free to pursue different internal distributive policies. (Note that this also does not preclude the possibility of redistribution at the interjurisdictional level.)

Second, I have not addressed incentives arguments against redistribution. However, it is important to stress that incentives arguments against redistributive *policies* are quite distinct from incentives arguments against distributive *moral positions*. Only the former are logically coherent. Regarding the latter, it makes no sense to say that

> what is untenable from the perspective of institutional robustness, therefore, is not the case for "common ownership" in the sense of a structure designed to fulfill a particular production or conservation goal, but the view that resources as a whole should be treated as the "common assets of humanity" with all people, by the mere fact of their existence, granted an equal right to a particular share. (143)

A distributive "view" in this sense is a moral principle, and it is a category mistake to think that a moral principle can be defeated by an empirical fact (Cohen 2008, ch. 5). A moral principle such as Rawls's difference principle defines a criterion of moral success (maximizing the goods of the least advantaged). One might say that *given* the empirical fact that redistribution reduces economic growth (let us suppose), the best way of fulfilling the difference principle is to enact a policy allowing people to claim private property in unowned land (for example). The empirical question would be whether this policy will help to fulfill the difference principle. No answer to this question, however, can affect the legitimacy or illegitimacy of the principle itself. Rejecting exaggerated versions of ideal theorizing does not entail rejecting the distinction between means and ends. Theories of justice aim to provide an end against which policies can be evaluated. A realistic theory of justice will make whatever concessions to reality are necessary to enable policies to be adopted that fulfill the posited end. Thus, the point of the difference principle is to allow for the fact that, due to incentive effects, maximizing the resources of the worst off may require allowing inequality. Rawls thus shields his theory from being

insensitive to empirical realities, even though no such reality can be adduced to defeat the justice of the difference principle.

Where does all this leave Pennington's argument for rejecting the egalitarian challenge to classical liberalism? As I see it, three roads are open to him, depending on where he wants to go.

First, it might be that Pennington's rejection of egalitarianism is not really driven by non-ideal considerations about knowledge and incentives, but instead by a substantial moral commitment to a libertarian conception of justice. Pennington's periodic emphasis on the necessity of unanimity points in this direction. For example, he writes that "people may wish to make them [trade-offs] differently, and when their values conflict it is doubtful that an agreed blueprint can be found to determine how much inequality is acceptable and how 'social' decisions should be made" (131). The assumption that decisions must be made unanimously points to a prior moral belief that each individual has a natural right to refuse to cooperate.

Another deep source of Pennington's rejection of distributive justice might be a thoroughgoing moral skepticism. An interpretation along these lines is suggested by the quoted passage on abortion, and by Pennington's use of scare quotes around the word *right*. If Pennington replied to my argument by simply denying that there is any such thing as justice or morality, this would certainly cut the legs out from under my critique. However, the fact that he argues that decentralization facilitates moral learning implies he believes there is something real to learn about. Moreover, if one were to adopt a radically skeptical stance, it is unclear to me on what grounds one could make any institutional recommendations at all, decentralized or otherwise. This, however, takes us into much deeper philosophical waters than I need to paddle in today.

The final road Pennington could take would be to abandon his resistance to egalitarianism and other theories of distributive justice. This would leave him able to continue raising incentives-based objections to redistributive *policies* insofar as they are supported by empirical evidence. It would place him in the company of "bleeding-heart libertarians" who combine hostility to the administrative or regulatory state with sympathy for the social-insurance or welfare state (Brennan 2012).

This road would be truer to the spirit of robust political economy: thinking about how the world can be improved given the deep importance of knowledge and incentive problems. At the very least, robust political economy may be agnostic about distributive justice, but it should not be atheist about it.

REFERENCES

Ackerman, Bruce, and James S. Fishkin. 2002. "Deliberation Day." *Journal of Political Philosophy* 10(2): 129–52.

Anderson, Elizabeth. 2006. "The Epistemology of Democracy." *Episteme* 3(1–2): 8–22.

Michael Bennett

Avi-Yonah, Reuven S. 2000. "Globalization, Tax Competition, and the Fiscal Crisis of the Welfare State." *Harvard Law Review* 113(7): 1573–1676.

Brennan, Jason. 2012. *Libertarianism: What Everyone Needs To Know*. Oxford: OUP.

———. 2014. "How Smart Is Democracy? You Can't Answer That Question a Priori." *Critical Review*: 1–26.

Caplan, Bryan. 2011. *The Myth of the Rational Voter: Why Democracies Choose Bad Policies*. Princeton University Press.

Cohen, G. A. 1994. "Back to Socialist Basics." *New Left Review* (207): 3–16.

———. 2008. *Rescuing Justice and Equality*. Cambridge, Mass: Harvard University Press.

DeCanio, Samuel. 2014. "Democracy, the Market, and the Logic of Social Choice." *American Journal of Political Science* 58(3): 637–52.

Dryzek, John S. 1996. *Democracy in Capitalist Times: Ideals, Limits, and Struggles*. New York: OUP USA.

Estlund, David. 2009. *Democratic Authority: A Philosophical Framework*. Princeton, N.J.: Princeton University Press.

Fearon, James D. 1998. "Deliberation as Discussion." In *Deliberative Democracy*, ed. Jon Elster. Cambridge: CUP.

Gerlsbeck, Felix. 2016. "What Is Democratic Reliability? Epistemic Theories of Democracy and the Problem of Reasonable Disagreement." *Critical Review of International Social and Political Philosophy*.

Guerrero, Alexander A. 2014. "Against Elections: The Lottocratic Alternative." *Philosophy & Public Affairs* 42(2): 135–78.

Hayek, F. A. 1949. "The Use of Knowledge in Society." In *Individualism and Economic Order*, Chicago: University of Chicago Press.

Keynes, John Maynard. 1933. "National Self-Sufficiency." *The Yale Review* 33(4): 755–69.

Knight, Jack, and James Johnson. 2011. *The Priority of Democracy: Political Consequences of Pragmatism*. New York: Princeton University Press.

Landemore, Hélène. 2013. *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton; Oxford: Princeton University Press.

López-Guerra, Claudio. 2011. "The Enfranchisement Lottery." *Politics, Philosophy & Economics* 10(2): 211–33.

Meade, J. E. 1964. *Efficiency, Equality and the Ownership of Property*. London: Allen & Unwin.

Miller, David. 2008. "Justice for Earthlings." In *Political Theory: Methods and Approaches*, eds. David Leopold and Marc Stears. Oxford ; New York: OUP Oxford.

Muldoon, Ryan. 2015. "Expanding the Justificatory Framework of Mill's Experiments in Living." *Utilitas* 27(2): 179–194.

Pennington, Mark. 2011. *Robust Political Economy: Classical Liberalism and the Future of Public Policy*. Cheltenham, UK ; Northampton, MA, USA: Edward Elgar Publishing Ltd.

Pincione, Guido, and Fernando R. Tesón. 2006. *Rational Choice and Democratic Deliberation: A Theory of Discourse Failure*. Cambridge; New York: Cambridge University Press.

Sen, Amartya. 2006. "What Do We Want from a Theory of Justice?" *The Journal of Philosophy* 103(5): 215–38.

Somin, Ilya. 2013. *Democracy and Political Ignorance: Why Smaller Government Is Smarter*. Stanford, California: Stanford University Press.

Stears, Marc. 2005. "The Vocation of Political Theory: Principles, Empirical Inquiry and the Politics of Opportunity." *European Journal of Political Theory* 4(4): 325–50.

Tiebout, Charles M. 1956. "A Pure Theory of Local Expenditures." *Journal of Political Economy* 64(5): 416–24.

Van Parijs, Philippe. 1997. *Real Freedom for All: What (If Anything) Can Justify Capitalism?* Oxford; New York: OUP Oxford.

Waldron, Jeremy. 1999. *Law and Disagreement*. Oxford; New York: OUP Oxford.