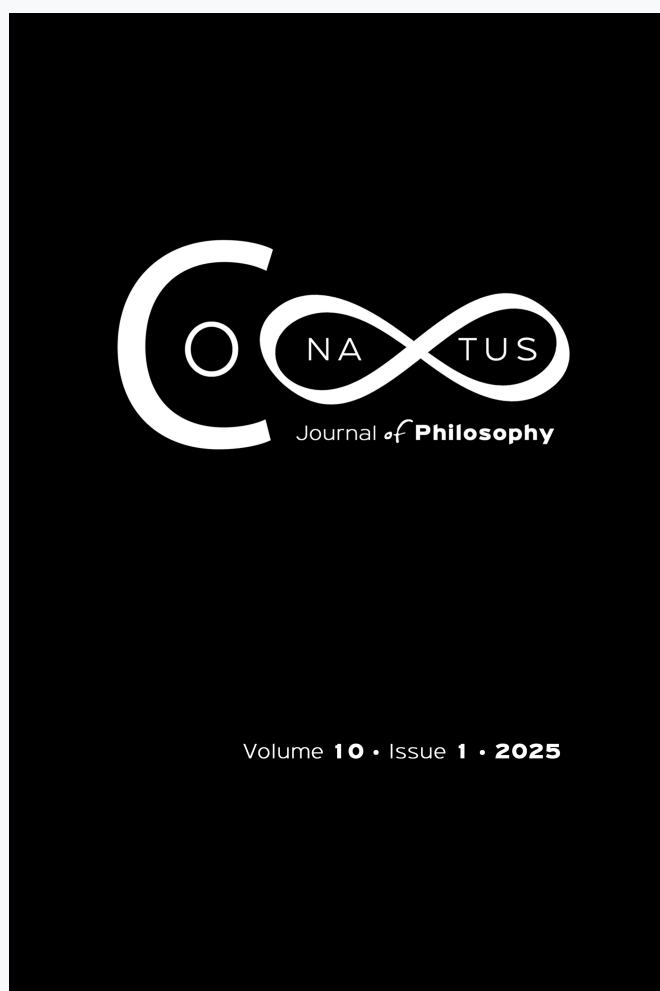


Conatus - Journal of Philosophy

Vol 10, No 1 (2025)

Conatus - Journal of Philosophy



Political Decision-making, Lottocracy, and AI

Joe Slater

doi: [10.12681/cjp.35215](https://doi.org/10.12681/cjp.35215)

Copyright © 2025, Joe Slater



This work is licensed under a [Creative Commons Attribution-NonCommercial 4.0](https://creativecommons.org/licenses/by-nc/4.0/).

To cite this article:

Slater, J. (2025). Political Decision-making, Lottocracy, and AI. *Conatus - Journal of Philosophy*, 10(1), 239–254.
<https://doi.org/10.12681/cjp.35215>

Political Decision-making, Lottocracy, and AI

Joe Slater

University of Glasgow, United Kingdom

E-mail address: joe.slater@glasgow.ac.uk

ORCID iD: <https://orcid.org/0000-0003-2269-7235>

Abstract

This article examines an argument for single-issue legislatures (SILs) as an alternative to typical decision-making procedures in representative democracies. It is argued that if this argument was successful, it could be extended to endorse decision-making processes utilizing advanced artificial intelligence. However, it is noted that this argument neglects an important feature of decision-making: authorization. It is important for its autonomy that a society must authorize certain decisions. If decisions were delegated to SILs or AI, this would undermine the autonomy of the group. This does not entail that these alternatives have no role in policy-crafting. It is argued that so long as a community can authorize a decision, perhaps by a vote, it need not undermine autonomy. An important caveat of this, however, is that the decision must be explicable to the community. For AI usage, this motivates the need for explainable AI (XAI).

Keywords: *lottocracy; AI; explanation and responsibility; authorization; autonomy*

I. Introduction

In representative democracies, elected officials are charged with crafting and enacting policies. Myriad problems have been long-documented for this system of policy-making. For example, given that the electorate is often uninformed about the issues, the politicians that get elected may be the ones who are most *persuasive* rather than the most competent. One suggestion, courtesy of Alexander Guerrero, is instead to use single issue legislatures selected by lot, which he terms “lottocracy.”¹ He argues persua-

¹ Alexander Guerrero, “Against Elections: The Lottocratic Alternative,” *Philosophy & Public Affairs* 42, no. 2 (2014): 135-178. Another lottocratic system is discussed in Claudio López-Guerra’s “The Enfranchisement Lottery,” *Politics, Philosophy & Economics*, 10, no. 2 (2010): 211-233.

sively that such a system may fare better than our current system in terms of its responsiveness to public opinion and good governance. However, if we explicitly express these as the values we want to promote, this may open the possibility of using sophisticated AI to develop policy instead, and this could be even more successful according to those metrics.

In this piece, I first consider Guerrero's argument for lottocracy, and show how (if we accept his assumptions), this argument could generalize and be used to defend policy-decisions being made by AI. I also note how similar defenses might be made for other existing (and proposed) uses of AI. Second, I propose that both of these systems fail to manifest something that is important about important decisions, namely that we *authorize* them. I point to several reasons why this authorization is valuable, drawing upon research in autonomy and responsibility. Finally, I investigate the implications of taking the value of authorization seriously. Crucially, I claim that we cannot authorize a decision unless we are sufficiently informed about it. This highlights the need for XAI in any domain where authorization is important.

II. Guerrero on lottocracy

Complaints about democracy are not new. In *The Republic* Plato argues for rule by "Philosopher Kings,"² viewing democracies as susceptible to the rule by the popular, rather than rule by the wise. In contemporary political philosophy, similar arguments are used by those endorsing "epistocracy" – the notion that experts should make the decisions.³ Alexander Guerrero has also argued that representative democracies have a variety of problems. Notably, he points to the possibility for politicians to be controlled by powerful interested parties. Because individual people cannot practically be fully informed on the range of complicated issues, people with power and resources can influence the political process in their favor, leading to policies that benefit them at the expense of others. This phenomenon – which he calls *capture* – undermines democracy's core principles of equality and fairness, as it allows a particular group to wield excessive political power and influence decision-making. It also makes it less likely that policies that

López-Guerra suggests that a system of universal suffrage is superior to the version he discusses. An extreme version of this, where a lottery decides one person who will make political decisions, is discussed in Isaac Asimov's short story, "Franchise" (1955). See Isaac Asimov, "Franchise," in *If: Worlds of Science Fiction*, ed. James L. Quinn, 2-15 (Quinn Publishing, 1955).

² Plato, *Republic*, trans. Robin Waterfield (Oxford University Press, 1993), V:473c.

³ Jason Brennan, *Against Democracy* (Oxford University Press, 2016).

benefit the electorate as a whole get enacted, as captured politicians will promote policy which favors those interest groups.

As an alternative, Guerrero proposes that legislation be determined by groups selected by lot – *lottocracy*. In Guerrero's suggestion, people would be randomly selected to a legislature about a particular domain of government, e.g., agriculture, tax policy or defense spending. These single-issue legislatures (SILs) would have about 300 people, demographically reflecting society as a whole, with legislators serving staggered three-year terms (each year, one hundred being replaced). Guerrero notes that this system would seriously reduce the risk of capture. In a representative democracy, paying a powerful politician (or someone who is expected to become a powerful politician) can be a good investment, as they could have a significant influence on decisions for a long time. Members of a SIL only have a small influence over one policy area, and for a mere three years, so if they could be influenced, this would be a less valuable investment for wealthy would-be manipulators.

The members of the SIL, Guerrero imagines, would meet several times a year, and be able to call experts to give advice. They would be well-compensated for their time, and encouraged to take their responsibility very seriously. This, Guerrero hopes, would enable discussions that would result in legislation that reflect the views of the general public, and would effectively pursue their policy goals, i.e., the system would be *responsive*. He also supposes that, freed from the interests of powerful parties, the resulting legislation would be better, however we might think that this could be understood from an objective vantage point, (e.g., may lead to fairer policies, or higher levels of average welfare), i.e., the system would result in *good governance*. In these two respects – responsiveness and good governance – Guerrero argues that lottocracy may perform better than representative democracy.

Of course, much can be said about whether Guerrero is correct about this claim. For instance, Lachlan Umbers argues that it is at best inconclusive whether these instrumental benefits would actually obtain. For instance, while members of the SIL might not be a suitable investment for capture, the experts who would be likely to provide advice would be, so wealthy people may provide extravagant support to fund favorable research. Furthermore, wealthy elites will retain lobbying power; they may be able to distort the information that SIL members receive, and SIL members may be more susceptible to such influences due to their relative inexperience in positions of political decision-making.⁴

⁴ Lachlan Montgomery Umbers, "Against Lottocracy," *European Journal of Political Theory* 20, no. 2 (2021): 312-334. Umbers also argues that lottocracies establish "objectionable social

For my purposes, I will assume that Guerrero is correct, i.e., that his version of lottocracy could yield more responsive policy decisions and improvements in improved governance. Even if this is the case, however, we might feel that there would be something important absent if we were to switch from representative democracy to lottocracy in our decision-making. Even though the laws arrived at might do a good job in terms of *reflecting* the values and goals of the population, the role of the general public is severely diminished. Christina Lafont touches upon this concern when she points out that lottocracy seems to endorse a “rule of the minority,”⁵ even if this minority is constituted by a representative sample and can accurately express public preferences. In the next section, I’ll suggest one element that I think is lacking in such a system, namely, the *public authorization* of the policies.

Before that, consider the following thought experiment. A team of psychologists, political theorists, computer scientists and philosophers team together to create an AI system to design legislation (either for a particular policy or set of policies). This AI – call it ‘Landru’⁶ – draws upon vast quantities of data from the whole population, perhaps via monitoring social media or surveys to representative samples (put aside for now questions about how to ethically collect such data or the logistics of doing so). Once Landru is operational, the policies it proposes are generally agreed to be far superior to the policies that our actual politicians arrive at. Perhaps this is because it is able to include in its calculations far more factors. After a while, politicians start to simply utilise policies proposed by Landru, rather than developing their own. Sometime later, the politicians decide (or perhaps Landru suggests!) getting rid of the politicians altogether, possibly because they have become glorified (and expensive) intermediaries. Why would anyone agree to this? I think there could be multiple motivations.⁷ Imagine a scenario where Landru’s recommended policy decisions have been taken for a long time. Eventually, a politician – to the surprise and outrage of the citizens – makes a decision contrary to the guidance, and this leads to some undesirable consequences. To

and political inequalities.”

⁵ Cristina Lafont, “A Militant Defence of Democracy: A Few Replies to my Critics,” *Philosophy and Social Criticism* 47, no. 1 (2021): 70.

⁶ This name is taken from the *Star Trek* episode “The Return of the Archons” (1967), where a society has allowed all policy to be determined by an advanced computer (Landru). It was also revisited in the *Star Trek: Lower Decks* episode, “No Small Parts” (2020).

⁷ I don’t think the details of *how* a society would arrive at such an arrangement are so important – what I’m really concerned with is whether the society that results would be deficient in some way.

avoid this, it might be demanded that faith be placed in Landru, which might be viewed as more trustworthy, or as a better judge of how to weigh the interests of competing groups equally. So a referendum is held on the future of policy-making, and a decision is reached to automatically defer to Landru. After this point, Landru makes all policy decisions. It continues to draw upon the wishes and values of the people in directing public policy.

If we were to assume that it was possible to develop an AI like Landru – I don't want to make grand claims about this here, but it at least seems like we can imagine developing an AI that would serve similar functions – how would this fare in terms of responsiveness and good governance?⁸ In terms of responsiveness, it seems like this system could be more effective than current representative democracies. This is because, perhaps by feedback mechanisms like regular surveys, or monitoring social media, it could detect changes in popular opinion and make policy changes at any point during the political cycle (not just when elections occur), could register viewpoints from every member of the population (and give them equal weight) and would be able to instigate policy changes immediately. It could also prove superior to the SILs in a lottocracy, because the members of the public summoned to service in the SIL would need quite a lot of time and effort to develop sufficient expertise about the policy areas. In terms of good governance, it also seems that Landru could perform very well here, however we think good governance should properly be understood. Once developed, it also looks like Landru would be a much cheaper option (and require less labor).

Now we are in a position to consider a generalization argument. Guerrero notes the importance of responsiveness and good governance. If (1) these features should determine our system of legislating,⁹ and (2) we developed a Landru-like AI (one that can meet the above conditions), then we ought to use it.

In what follows, I will assume that it is possible for us to develop AI like this. I take this strategy because I want to demonstrate the

⁸ In this thought experiment, I simply stipulate that programming Landru to make ethical decisions (e.g., decisions that do not violate rights, accompanied by no desires to dominate humanity). For some discussion about the difficulties in creating an AI of this sort, see Michael Anderson et al., "Towards Moral Machines: A Discussion with Michael Anderson and Susan Leigh Anderson," *Conatus - Journal of Philosophy* 6, no. 1 (2021): 177-202.

⁹ Guerrero does not endorse this claim, which is something of an oversimplification, but one I take to be worth examining. Rather, he says that these two factors present serious problems for electoral democracy, and that he thinks lottocracy is normatively attractive because it is superior in these dimensions.

failings of the first condition. When the failure of the first condition is revealed, this cuts off the inference from the potential instrumental successes of lottocracy (or Landru-cracy) to the claim that these systems should be implemented. Furthermore, observing how this argument fails may highlight issues with other AI uses, which perhaps we should be suspicious of, for instance, the use of algorithms in sentencing, or evaluation of résumés.

III. The Importance of authorization

In both of the alternatives to representative democracy considered in the previous section, the role of the public in decision-making has been diminished. Importantly, if there is a problem with these systems, it isn't that bad policies are being selected (by stipulation), or that the legislation passed is not what the public would have selected. Each system is designed to result in legislation that the public would be in favor of, if ideal conditions obtained (e.g., that they were fully informed about the issues, had time to deliberate, and had no cognitive limitations in foreseeing consequences). There are still various different ways we might view these systems as lacking. On one hand they could be seen as deficient in terms of accountability.¹⁰ Such a system may also have undesirable consequences with regard to the public becoming disengaged with political matters.¹¹ A further missing feature, I argue, is *authorization*.

To illustrate what I mean by “authorization,” consider someone employed as an assistant. There may be standing duties that are always part of this job. Some of these may not be fully determined in a variety

¹⁰ Hubertus Buchstein, “Lottocracy and Deliberative Accountability,” *Philosophy and Social Criticism*, 47, no. 1 (2021): 40-44; see also Kate Crawford and Jason Schultz, “AI Systems as State Actors,” *Columbia Law Review* 119, no. 7 (2019): 1941-1972. They discuss the ‘accountability gap,’ which currently haunts algorithmic decision-making. See Michael Townsen Hicks, James Humphries, and Joe Slater, “ChatGPT is Bullshit,” *Ethics and Information Technology* 26, no. 38 (2024). Hicks et al. raise the problem of accountability for the falsehoods uttered by large language models. See also Joe Slater, James Humphries, and Michael Townsen Hicks, “ChatGPT Isn’t ‘Hallucinating’ – It’s Bullshitting,” *Scientific American* 34, no. 1 (2025): 46-47.

¹¹ Cristina Lafont, *Democracy without Shortcuts: A Participatory Conception of Deliberative Democracy* (Oxford University Press, 2020). Because no such system exists, and the notion has not inspired considerable debate, there is not a wealth of discussion about disengagement as a result of AI use in policy decisions. However, the reasons this may occur are the same as with lottocratic systems, namely, that if individuals see themselves as having no direct role in decision-making, they may lack the motivation to remain informed on political issues. This is essentially the counterpart of an argument made by Mill for democracy, namely that an empowered citizen is “called upon ... to weigh interests not his own; to be guided, in cases of conflicting claims, by another rule than his private partialities.” John Stuart Mill, *Considerations on Representative Government* (Parker, Son, and Bourn, 1861).

of respects; i.e., they permit some discretion. For example, if not specified, documents might be formatted in a variety of ways, or lunch could be purchased at different venues. But there are also likely to be a variety of tasks that an employer will request to be done, e.g., to arrange a meeting or order new furniture. If the assistant was to perform these tasks without the employer's request (or being told that they could do them), the employer could appropriately feel wronged by this, even if the employer would have made these requests later anyway.

To embellish the example further, and make it more similar to a legislative decision-making process that citizens have voted for, we could imagine that the employer has instructed their assistant that they can make decisions on a wide range of matters, including important decisions normally seen within the employer's remit. In this case, when the secretary makes the decisions, they have not wronged the employer. However, the employer might still feel, when these decisions are enacted, like something is missing from their involvement. They haven't made the decision themselves (even if they would have made the same decision).

Why should it matter for political decision-making whether the citizens *brought about* the decisions, as opposed to the decisions merely reflecting what they would opt for when fully informed? I suggest that the answer lies in autonomy. In ethical theory, autonomy is seen as having central importance by philosophers of varying different stripes, as seen in Kant,¹² Mill,¹³ or Rawls.¹⁴ Mill sees respecting autonomy as conferring certain benefits, such as promoting utility. Others, like Kant, regard autonomy as intrinsically valuable; an individual's being autonomous confers upon them a special moral status, which commands respect.¹⁵ Regardless of why we see it as valuable, *that* we value autonomy is uncontroversial.¹⁶

What *is* controversial is exactly how we should understand autonomy. There are three broad types of views of autonomy: reason-based

¹² For a discussion of Kant's conception of autonomy see Janis Schaab, "Kant on Autonomy of the Will," in *The Routledge Handbook of Autonomy*, ed. Ben Colburn, 44-54 (Routledge, 2023).

¹³ This is especially clear in John Stuart Mill, *On Liberty* (Broadview Press, 1859).

¹⁴ For example, Rawls notes that his first principle of justice can be given an interpretation based on Kant's notion of autonomy. See John Rawls, *A Theory of Justice, Revised Version* (Harvard University Press, 1999), 221.

¹⁵ For a discussion of these views see Mark Piper, "Justifying Respect for Autonomy," in *The Routledge Handbook of Autonomy*, ed. Ben Colburn, 293-302 (Routledge, 2023): 293-302.

¹⁶ Even Sarah Conly, who argues that we place too much value on autonomy, accepts that we do value autonomy. See Sarah Conly, *Against Autonomy* (Cambridge University Press, 2014).

views, motivation-based views, and self-creation/self-authorship views.¹⁷ Notorious among the reason-based views is the Kantian picture. This holds that the notion of self-legislation is important: with the use of reason, one can recognize the unconditional authority of the moral law. The rational agent will then make themselves act only on maxims that are consistent with the moral law, expressed by the categorical imperative.

On the motivational view, expressed by Gerald Dworkin, the structure of one's motivations was emphasized. Autonomy on such a view involves having authentic preferences that do not have some alien or external source.¹⁸ Preferences being authentic requires that they are those of one's "true self" or "highest-order self," i.e., not only what one wants, but what one *wants to want*.

Self-creation views were advocated by Joseph Raz. These highlight the importance of the person *determining* how their life goes. Raz sees autonomy as "an ideal of self-creation or self-authorship."¹⁹

Strictly speaking, the positions mentioned in the aforementioned paragraphs are views of *personal* autonomy. However, we also talk about *group* autonomy. We may ask, for instance, if some policy violates the autonomy of an indigenous group, or whether some treaty diminishes the autonomy of a nation. As Wellman notes: "Group autonomy exists when the group as a whole, rather than the individuals within the group, stands in the privileged position of dominion over the affairs of the group."²⁰

Just as we regard considerations of *personal* autonomy as providing us with reasons for or against some decision, so too do we with regard to these kinds of groups.²¹ It is reasonable to wonder how well views of personal autonomy could map onto accounts of *groups*. This is particularly evident when we consider puzzles of group agency. Can a group *really* be an agent? Does it make sense to say that a group *believes* X, or *wants to do* Y? This is a rich debate, which I cannot do justice to here.²² What is clear is that we at least talk *as if* groups can have

¹⁷ Ben Colburn, *Autonomy and Liberalism* (Routledge, 2010).

¹⁸ Gerald Dworkin, "Autonomy and Behavior Control," *Hastings Center Report* 6, no. 1 (1976): 23-28.

¹⁹ Joseph Raz, *The Morality of Freedom* (Oxford University Press, 1988), 369.

²⁰ Christopher Heath Wellman, "The Paradox of Group Autonomy," in *Autonomy. Social Philosophy and Policy*, eds. Ellen Franken Paul, Fred D. Miller Jr, and Jeffrey Paul, 265-285 (Cambridge University Press, 2003), 273.

²¹ Exactly *which* kinds of groups are such that their autonomy is a morally relevant consideration is a difficulty question.

²² For an overview of some of the discussions in this area, see Katherine Ritchie, "The Meta-

such states and can act as groups. For instance, we may say, perfectly felicitously, that the UK voted the Tories out in the 2024 general election. We do, in everyday discourse, speak as if groups are agents that can rationally consider options, make decisions, and act accordingly.²³

Given that we can talk about groups having the same kinds of properties as are required in the accounts of personal autonomy, then – if we think that group autonomy is simply the group equivalent of personal autonomy – it seems natural to think of group versions of the accounts of personal autonomy, i.e., that we can talk about reason-based views, motivation-based views and self-creation/self-authorship views of group autonomy. So, for my purposes, I assume that similar conditions will be applicable.

On each of these accounts of autonomy (or at least, on certain versions of these accounts), I contend that a persuasive story can be told about the value of authorization with regards to important decisions.²⁴ When we consider *personal* autonomy, it is important that we authorize certain decisions about how our own lives go. When we consider the autonomy of a *group* or *nation*, similarly, the group must authorize certain decisions about what affects them. In what follows, I will pay particular attention to self-authorship views, which are particularly suitable for such an explanation, and specifically examine Ben Colburn's.²⁵

In Colburn's view, "autonomy requires both deciding for oneself what is valuable and also living in accordance with that decision by making decisions that make one responsible for the shape one's life takes."²⁶ An important part of this view is *making your life go the way you decide*. This is something that we care about, and perhaps see as part of us living a good life. To be autonomous in this way, we do need to be *responsible* for the shape of our life; we need to actually bring it about.

For individuals or groups to be autonomous, they must actually *be* the authors of their lives. The autonomy critique of lottocracy/Landru-crazy, however, holds that there is simply a problem because the citizens themselves have forfeited important controls over their lives.

physics of Social Groups," *Philosophy Compass* 10, no. 5 (2015): 310-321.

²³ For a discussion of how group intention is possible, see Margaret Gilbert, "Shared Intentions and Personal Intentions," *Philosophical Studies* 144, no. 1 (2009): 167-187.

²⁴ Colburn himself makes a similar move in discussions of assisted dying, in Ben Colburn, "Autonomy, Voluntariness and Assisted Dying," *Journal of Medical Ethics* 46, no. 5 (2020): 316.

²⁵ Colburn's view is described in detail in *Autonomy and Liberalism*.

²⁶ Colburn, "Autonomy," 316.

They are no longer masters of their lives, forging their own destinies.

It is perhaps unsurprising that these alternative systems fail to manifest conditions of autonomy. Autonomy is often characterized as self-rule or self-governance, and this is precisely what is sacrificed when decisions are entirely delegated to a randomly selected lot, or to an AI. To put it in terms of Lincoln's proverb, the governing may still be *for* the people, but in an important sense, it is no longer *by* the people.

When we attend to considerations of autonomy, this can elucidate issues that may have been obscured in various analyses of policy-making systems. This can be seen in the case of a pure lottocracy, where policy decisions were all determined by randomly selected legislators. Guerrero suggests that this could result in policies more responsive to citizens' values and objectively better at whatever we would desire government to do. But as these policy decisions are important to our lives – perhaps to our very identities – our being excluded from participation in the decision-making process constitutes a violation of our autonomy. And if autonomy is one of our fundamental values, this would be a bitter pill to swallow.

Thus far, I have argued that moving from a democracy – either direct or representative (voting for representatives, as well as for particular policy proposals can be regarded as a form of authorization) – to a system where laws are designed by a representative sample of the population or by a sophisticated AI fails to respect the value of autonomy. Finally, I reflect upon what the requirement for authorization really entails, and what implications this has for our policy-making decisions and other areas we might utilize AI.

IV. Authorization, AI and XAI

If we accept, as I have suggested we should, that it is important for our autonomy that we authorize important decisions, this leaves open several questions. For instance, we might wonder what form the authorization must take, how often it must occur, or how fine-grained acts of authorization must be. If explicit authorization for every aspect of every policy must occur, then seemingly no system but a direct democracy with compulsory voting would suffice. On the other end of the spectrum, we might imagine that our imagined society does have a vote every fifty years to decide whether or not to continue allowing Landru to determine all policy decisions. If something like this – voting to continue with the AI-system once a generation (or an alternative whereby citizens voted whether or not to continue a lottocratic system at similar intervals) – would constitute sufficient authorization then

rule by lot or AI would presumably be acceptable, so long as these kinds of “approval elections” did take place sufficiently regularly.

I do not propose to give a precise answer to such questions here, but will note a couple of important features that the authorization should have.

First, while it is possible to talk about a certain *decision* being made autonomously, autonomy is typically viewed as a *global* property, i.e., it applies over a long period of time, not merely with regards to particular decisions.²⁷ It does admit of degrees, but above a certain threshold we will attribute the property. So to have this property, the entity in question (be it a person or a group) must have enough decisions of enough importance. With this in mind, for public to act autonomously in deciding its policy goals, its role cannot be a one-time affair. It must persist over a course of time, e.g., with regular votes. This may also allow for systems where citizens can participate in a variety of other ways – I take no stand on this here.

Second, when we consider how fine-grained an instance of authorization must be in order, it could again be helpful to consider the value we are trying to promote. Recall the claim that autonomy requires, as Colburn puts it, “deciding for oneself what is valuable and also living in accordance with that decision by making decisions that make one responsible for the shape one’s life takes.” Highlighting *responsibility* here can be fruitful. When Colburn discusses this,²⁸ he notes two senses we can be responsible: attributability and substantive responsibility. Colburn sees both of these as necessary for the required responsibility.

For our actions to be attributable to us, “it must be recognizably our choices and actions that make our lives the way they are.”²⁹ To be substantively responsible, we must be “liable for the consequences of the things attributable to us.”³⁰ Exactly what this responsibility requires may be very context-dependent. In some contexts, the public indicating a general direction to go on might be all that is necessary for them to meet these conditions. In others, perhaps something more specific could be needed. What does seem important is that the public has been *informed*. To be liable for consequences of decisions, clearly one needs to understand (or have *the ability* to understand) at least something about the content of those decisions. Again, being informed might re-

²⁷ Ibid., 4.

²⁸ When discussing responsibility, Colburn draws upon a discussion from Thomas M. Scanlon’s book. See Thomas M. Scanlon, *What We Owe to Each Other* (Harvard University Press, 1998), 148-251.

²⁹ Colburn, “Autonomy,” 32.

³⁰ Ibid., 32.

quire different degrees of detail in different scenarios, but complete ignorance about policies selected and reasons that count in favor of them seems incompatible with this.

Now we are in a position to say something about what authorization must look like. Because autonomy is a global condition, authorization must be regular. The public must have more than merely occasional say about policy-matters. Furthermore, in order for the policies to be properly attributable to the public, they must be sufficiently *informed*.

The alternative decision-making systems can now be considered again. Lottocracy does not respect autonomy, because the public does not authorize the decisions. However, if the lottocratic method of crafting legislation is combined with suitable public participation to ensure authorization, this problem can be overcome. Members of the legislator can explain their decisions to members of the public, and then a vote can be used to ratify the decision. Interestingly, this is compatible with Buchstein's proposal,³¹ and with various ways citizens' assemblies have been used recently, whereby the assemblies determine what questions in a referendum should be put to the public.³²

With regards to legislation crafted by AI, a problem becomes apparent. One worry with machine learning algorithms is that they essentially constitute "black boxes,"³³ i.e., we may know what goes in and what comes out, but the process is left something of a mystery to us. If Landru, our policy-making AI, was of this sort, then it would not be possible for us to authorize the decisions it arrives at, because the public could not be informed.

However, just as the lottocracy case could be adjusted, so can the AI case. The public would need some say on the policies proposed. But more than that, they would need access to the *reasons* why these policies are favored over alternatives. This would require explainable AI (XAI), i.e., "systems that explain how the algorithms reach their conclusions or predictions."³⁴ This explanation would need to be suitable for

³¹ Buchstein, "Lottocracy."

³² For example, Ireland's use of a citizens' assembly, which set the question that would be put to a referendum on abortion. See "The Irish Abortion Referendum: How a Citizens' Assembly Helped to Break Years of Political Deadlock," *Electoral Reform Society*, accessed June 9, 2019, <https://www.electoral-reform.org.uk/the-irish-abortion-referendum-how-a-citizens-assembly-helped-to-break-years-of-political-deadlock/>.

³³ For some discussion of "black box" issues with algorithms see Avihay Dorfman and Alon Harel, "Why Not Artificial Intelligence?" in *Reclaiming the Public*, ed. Avihay Dorfman and Alon Harel (Cambridge University Press, 2024), 171-187.

³⁴ Ashley Deeks, "The Judicial Demand for Explainable Artificial Intelligence," *Columbia Law*

at least some human beings to understand. It need not be put in terms so simple that every member of the public can understand (at least not in principle), as some members of the public will be able to take on the task of explaining to others.

Worries about use of algorithms has led to the demand for XAI in other areas. For instance, algorithms used in sentencing criminals,³⁵ to decide outcomes of loan applications or to determine which students are admitted into universities.³⁶ Often, the concern with such uses is that the algorithms are giving weight to factors that should not be relevant, and if the algorithm is a black box, parties disadvantaged by this have no grounds for complaint. For example, companies have used algorithms to evaluate CVs of candidates, but some of these algorithms have systematically disadvantaged women, or paid undue attention to a candidate's name.³⁷ These issues are concerning even putting aside considerations of autonomy of the organizations using them, and has led to calls for a right to protest decisions made by AI.³⁸ Yet even if the algorithms did not exhibit these significant failings with respect to fairness, if organizations value their autonomy, they may see these algorithms as independently objectionable. If no individual members of the organization – be it an employer making a hiring decision, a university considering who to admit, or a court deciding on sentencing – have knowingly authorized the decision, it is dubious to what extent the decision can appropriately be attributed to the organization.

V. Conclusion

I have argued that, when we consider political decision-making, focusing on features like responsiveness and efficiency in obtaining results

Review 119, no. 7 (2019): 1829.

³⁵ Some of the issues with the use of algorithms in sentencing are discussed by Duncan Purves and Jeremy Davis. See Duncan Purves and Jeremy Davis, "Should Algorithms that Predict Recidivism Have Access to Race?" *American Philosophical Quarterly* 60, no. 2 (2023): 205-220.

³⁶ Gyorgy Denes, "A Case Study of Using AI for General Certificate of Secondary Education (GCSE) Grade Prediction in a Selective Independent School in England," *Computers and Education: Artificial Intelligence* 4 (2023): 100-129.

³⁷ "Companies Are on the Hook if Their Hiring Algorithms are Biased," *Quartz*, accessed August 13, 2023, <https://qz.com/1427621/companies-are-on-the-hook-if-their-hiring-algorithms-are-biased>.

³⁸ Problems with algorithm use in hiring are also discussed by Pauline T. Kim and Matthew T. Bodie. See Pauline T. Kim and Matthew T. Bodie, "Artificial Intelligence and the Challenges of Workplace Discrimination and Privacy," *ABA Journal of Labor and Employment Law* 289, no. 35 (2021): 289-315. Margot Kamiski and Jennifer Urban, "The Right to Contest AI," *Columbia Law Review* 121, no. 7 (2021): 1957-2048.

can obscure some of what we value. Specifically, this approach fails to appreciate the value we place on our decisions being *autonomous*. We want to be the authors of our own lives; we hold in high esteem our “privileged position of moral dominion”³⁹ over our affairs, and see alternative ways of life which deprive us of this as problematic. This is highlighted by an unease some of us feel regarding lottocracy. It can also be seen in the argument I constructed about a policy-making AI. If we accept that a Landru- like AI would create good policy decisions, which are extremely responsive to public values, we can arrive at a disjunctive conclusion. Either i) AI-legislation (at least regarding some domains of decisions) is superior than both democratic *and* lottocratic methods, or ii) the values of political decisions are not exhausted by the two main factors Guerrero considers. I argue in favor of ii), suggesting that there is value in the *authorization* of a decision. This feature is lacking in lottocratic and AI-crafted legislation. The lack of authorization inhibits the public’s autonomy.

In order for a society’s decisions to be made autonomously, it must authorize those decisions. This entails that the public – as a whole, not merely a small subset – has a say (e.g., elections), and that the justifications for the decisions are made clear. The easiest ways this can be done, utilising new methods of legislation-crafting, would still involve elections, perhaps with the matters put to an electorate determined not merely by politicians. Those excited by the potential for new technological developments to improve our decision-making (and, as a result, our lives) are right to be excited, but if we value our autonomy, we should be careful not to sacrifice that along the way.

References

Anderson, Michael, Susan Leigh Anderson, Alkis Gounaris, and George Kosteletos. “Towards Moral Machines: A Discussion with Michael Anderson and Susan Leigh Anderson.” *Conatus – Journal of Philosophy* 6, no. 1 (2021): 177-202.

Asimov, Isaac. “Franchise.” In *If: Worlds of Science Fiction*, edited by James L. Quinn, 2-15. Quinn Publishing, 1955.

Brennan, Jason. *Against Democracy*. Oxford University Press, 2016.

Buchstein, Hubertus. “Lottocracy and Deliberative Accountability.” *Philosophy and Social Criticism* 47, no. 1 (2021): 40-44.

³⁹ Wellman, “The Paradox of Group Autonomy,” 263.

Colburn, Ben. *Autonomy and Liberalism*. Routledge, 2010.

Colburn, Ben. "Autonomy, Voluntariness and Assisted Dying." *Journal of Medicine and Philosophy* 46, no. 5 (2020): 316-319.

Conly, Sarah. *Against Autonomy*. Cambridge University Press, 2013.

Crawford, Kate, and Jason Schultz. "AI Systems as State Actors." *Columbia Law Review* 119, no. 7 (2019): 1941-1972.

Deeks, Ashley. "The Judicial Demand for Explainable Artificial Intelligence." *Columbia Law Review* 119, no. 7 (2019): 1829-1850.

Denes, Gyorgy. "A Case Study of Using AI for General Certificate of Secondary Education (GCSE) Grade Prediction in a Selective Independent School in England." *Computers and Education: Artificial Intelligence* 4 (2023): 100-129.

Dorfman, Avihay, and Alon Harel. *Reclaiming the Public*. Cambridge University Press, 2024.

Dworkin, Gerald. "Autonomy and Behavior Control." *Hastings Center Report* 6, no. 1 (1976): 23-28.

Gershgorn, Dave. "Companies Are on the Hook if Their Hiring Algorithms Are Biased." *Quartz*, October 22, 2018. <https://qz.com/1427621/companies-are-on-the-hook-if-their-hiring-algorithms-are-biased>.

Gilbert, Margaret. "Shared Intentions and Personal Intentions." *Philosophical Studies* 144, no. 1 (2009): 167-187.

Guerrero, Alexander. "Against Elections: The Lottocratic Alternative." *Philosophy & Public Affairs* 42, no. 2 (2014): 135-178.

Hicks, Michael Townsen, James Humphries, and Joe Slater. "ChatGPT is Bullshit." *Ethics and Information Technology* 26, no. 38 (2024): 1-10.

Kamiski, Margot, and Jennifer Urban. "The Right to Contest AI." *Columbia Law Review* 121, no. 7 (2021): 1957-2048.

Kant, Immanuel. *Groundwork of the Metaphysics of Morals. A German-English Edition*. Translated by Mary Gregor and Jens Timmermann. Cambridge University Press, 2023.

Kim, Pauline T., and Matthew T. Bodie. "Artificial Intelligence and the Challenges of Workplace Discrimination and Privacy." *ABA Journal of Labor and Employment Law* 289, no. 35 (2021): 289-315.

Lafont, Cristina. *Democracy without Shortcuts: A Participatory Conception of Deliberative Democracy*. Oxford University Press, 2020.

Lafont, Cristina. "A Militant Defence of Democracy: A Few Replies to my Critics." *Philosophy and Social Criticism* 47, no. 1 (2021): 69-82.

López-Guerra, Claudio. "The Enfranchisement Lottery." *Politics, Philosophy & Economics* 10, no. 2 (2010): 211-233.

Macleod, Christopher. "Mill on Autonomy." In *The Routledge Handbook of Autonomy*, edited by Ben Colburn, 75-84. Routledge, 2023.

Mill, John Stuart. *On Liberty*. Broadview Press, 1859.

Mill, John Stuart. *Considerations on Representative Government*. Parker, Son, and Bourn, 1861.

Oshana, Marina. *Personal Autonomy in Society*. Routledge, 2006.

Piper, Mark. "Justifying Respect for Autonomy." In *The Routledge Handbook of Autonomy*, edited by Ben Colburn, 293-302. Routledge, 2023.

Plato. *Republic*. Translated by Robin Waterfield. Oxford University Press, 1993.

Purves, Duncan, and Davies Jeremy. "Should Algorithms that Predict Recidivism Have Access to Race?" *American Philosophical Quarterly* 60, no. 2 (2023): 205-220.

Rawls, John. *A Theory of Justice: Revised Edition*. Harvard University Press, 1999.

Raz, Joseph. *The Morality of Freedom*. Oxford University Press, 1988.

Ritchie, Katherine. "The Metaphysics of Social Groups." *Philosophy Compass* 10, no. 5 (2015): 310-321.

Scanlon, Thomas M. *What We Owe To Each Other*. Harvard University Press, 1998.

Schaab, Janis. "Kant on Autonomy of the Will." In *The Routledge Handbook of Autonomy*, edited by Ben Colburn, 44-54. Routledge, 2023.

Slater, Joe, James Humphries, and Michael Townsen Hicks. "ChatGPT Isn't 'Hallucinating' – It's Bullshitting." *Scientific American* 34, no. 1 (2025): 46-47.

Umbers, Lachlan Montgomery. "Against Lottocracy." *European Journal of Political Theory* 20, no. 2 (2021): 312-334.

Wellman, Christopher Heath. "The Paradox of Group Autonomy." In *Autonomy. Social Philosophy and Policy*, edited by Ellen Franken Paul, Fred D. Miller Jr, and Jeffrey Paul, 265-285. Cambridge University Press, 2003.