LLM representatives, autonomy, and means-end trades

Abstract: Manuali (2026) argues that LLM representatives are incompatible with democratic autonomy. I argue to the contrary. The key point is that if we create and empower LLM representatives, we trade influence over our choice of means for influence over our ends. As we care more about the latter than the former, this augments our autonomy overall.

Introduction

Imagine having an LLM-based agent that acts as your representative in political affairs. You discuss your values and policy views with this agent. It accurately infers what you would want in areas you haven't considered in any detail. It asks you if it is unsure about your views on some important issue. It checks whether it has accurately identified your values and it faithfully transmits them into action. You could, today, use this agent to contact office-holders to convey your views. You could use it to write letters to your representative. More radically, we might empower it to vote in your stead. Think of elections or referendums where you know little about the candidates or issues. We could allow you to (revocably) delegate your vote to your LLM representative in these contexts. Several people have thought that there would be some democratic value to creating and empowering LLM representatives of this sort (Gudiño-Rosero et al. 2024; Lavi 2025; Jarrett et al. 2025; Han 2025, 16–17). The idea is that they would extend each person's influence into areas in which they are currently impotent.

Lorenzo Manuali (2026) has recently argued against such LLM representatives. He thinks empowering them would threaten democratic autonomy. This rests on the thought that democracy is justified, in part, due to its connection to such autonomy. The idea here is that, just as it is valuable for you to be the author of your personal affairs, so it is valuable for you to be the author of your social and political affairs (Lovett and Zuehl 2022; Lovett 2020, 2026a, 2026b). Such authorship requires that your intentions causally contribute to what the government does; democracy facilitates such a connection. The problem with LLM representatives, for Manuali, is that they do not engender this causal connection. LLM representatives predict what you would want on issues you have never considered. You likely don't have any intentions concerning these issues. Even if you do, they aren't what cause the LLM representative to predict that you would want a certain policy option. So, when such a representative effectively

pushes for that option, there is no causal connection between your intentions concerning the policy and that policy being enacted. So you enjoy no authorship of the policy. Manuali is thus pessimistic about the use of LLM representatives.

I am more optimistic. In this commentary, I explain why.

1. Dubious Views

What Manuali gets right, in my view, is that there are some unpromising views around about how to *justify* the creation of such representatives. Some people seem to think that the democratic value of LLM representatives is that they would improve the match between government policy and people's preferences (Gudiño-Rosero et al. 2024). This is of dubious value. You are not the author of government policy just because the policy matches your preferences. You might agree with everything the dictator does. But they don't care about your opinion; the agreement is accidental. Here, you don't author government policy. Some people, instead, seem to think that the value of LLM representatives is that they would improve the match between government policy and what our preferences would be were we better informed (cf. Konya et al. 2023). Again, this is of dubious value. Suppose your parents wisely arranged your life so you would endorse your marriage partner were you better informed, but you actually detest them. Here you are not the autonomous author of who you marry. For that, your actual intentions must causally connect to your romantic outcomes. Yet connecting political outcomes to an intention you would have had in hypothetical circumstances does not connect them to any intention you actually do have. So these are not good ways of understanding the democratic value of LLM representatives.

2. Means-End Trades

There is, however, a different way to understand the potential value of LLM representatives. We can think of it in terms of a tradeoff between influence over your means and influence over your ultimate ends. Consider what happens when you hire an expert to do some job. Suppose you're building a house, but you know little about electrical wiring. You hire a good electrician to do the wiring. You just want the lights to turn on without burning the house down. So you tell the electrician to use their judgement to this end. You get a functional house with safe wiring. Here you don't author the exact placement or color of the wiring in your house. Yet you are the author of the ultimate outcome: that your house has good wiring. You told the electrician your ultimate end, and through their work you achieved it. You could have influenced the details of the electrician's work—by peppering them with questions, offering your advice—but this would likely have worsened the ultimate outcome. So in this case, you increase your authorship over

your ultimate end by sacrificing some authorship over the means to that end. That seems to improve your autonomy overall—it seems to increase your opportunity for authoring your affairs.

To understand this theoretically, we need an account of overall autonomy. To evaluate how autonomous someone is, we look at how much they can causally contribute to a range of outcomes (Lovett 2026a). The first key point is that the relevant contributions are contributions you know how to make (cf. Lovett 2020). There is a sense in which I am able to install safe wiring in my house myself. I can physically do the actions. But I don't know which actions these are. To have an opportunity for authoring this outcome in the relevant sense, I must know which of my actions contribute to it. Autonomy requires knowing how to contribute to different outcomes. This is why delegating choice of means to experts can increase one's opportunity for authoring an end. An expert knows more than you about which are the effective means to your ends. So letting them choose the means can increase the extent to which you knowingly contribute to the end.

The second key point is that not all outcomes matter equally to autonomy. You care about some things more than others. Influence over outcomes you care more about contributes more to your autonomy (Lovett 2026a, 7–8). We can see this in our personal life. The extent to which you can author who you marry or what career you pursue probably matters more to your overall autonomy than the extent to which you can author what brand sneakers you wear. This is so when, and because, you care more about the former than the latter. The point extends to political outcomes. If you don't really care who is state mine inspector in Arizona or whether proposition 45 passes, authorship of these things won't much contribute to your autonomy. If you care much more about protecting the environment or promoting economic growth, authorship over these latter outcomes matters much more to your autonomy. Thus, if you can trade influence over something you care little about for influence over something you care more about, this will increase your autonomy overall. More generally, we can think of overall autonomy as a measure of how much one can knowingly contribute to each of a range of outcomes weighted by how much one cares about the outcome.

Let's bring these points together. Observe that we see some outcomes as mere means to our ends. That just means we care about some outcomes primarily because they causally contribute to a further outcome. In these cases there are often multiple means to our ends. So we tend to care less about utilizing any particular means than we do about achieving our ultimate end. In some such cases, experts are better at evaluating the efficacy of each possible means than we are. Thus, delegating to such an expert the choice of means gives us more influence over whether the end occurs. This is a means-end trade. It gives you more influence over an outcome you care more about—the achievement of the end—

at the cost of less influence over outcomes you care less about—the choice of means. The influence over the ends even flows from your intention to bring them, specifically, about. If influence over things you care more about contributes more to your autonomy, this trade augments your autonomy overall. Delegating the choices of means to experts broadens your opportunity to author your affairs.

I myself became interested in such means-end trades by considering public bureaucracies. If we know anything about providing effective, efficient public services, it is that it requires a meritocratic civil service (Gallo and Lewis 2012; Xu 2018). We must hire and promote people on ability, rather than on political loyalty. Meritocracy means you get better people in each role. It incentivizes all public officials to work harder. But establishing a meritocratic civil service means citizens lose influence over some decisions. They lose authorship of who exactly fills certain roles in the public bureaucracies. They likely lose influence over the fine-grained decisions these people make. Is this antithetical to democratic autonomy? I think not. By establishing a meritocratic bureaucracy, we become more able to provide government services effectively. We become better able to build infrastructure, provide schooling, disburse benefits. We care more about these things than the means to them. So this trade improves our opportunity for authoring our social and political affairs overall. Meritocratic civil services sacrifice influence over means for influence over ends. To capture the worth of this, we must, I think, accept that means-end trades contribute to autonomy.

3. LLM Representatives

Let's go back to LLM representatives. Here, an important background fact is that ordinary voters don't know much about politics. We don't know who is running in each election, what their policies are and what the effects of those policies might be. As Philip Converse famously put it, "the pithiest truth I have achieved about electorates is that where political information is concerned, the mean level is very low…" (Converse 2000, 331). Pertinently, when it comes to petitioning officials, voting in elections, or casting a ballot in a referendum, we often don't know what actions will serve as effective means to our ends. So our influence over the outcome of such processes does not much increase our opportunity to knowingly contribute to such ends. Such conditions make delegating our vote to a more knowledgeable proxy an especially effective way to increase our opportunity to influence these ends. Specifically, if we empower an LLM representative to act on our behalf, we may lose influence over exactly how many votes a specific candidate wins or each ballot measure receives. But we increase our influence over what happens in our society writ large. If we care more about the latter more than the former, this is a winning means-end trade. It increases our overall autonomy.

Let us see how this works in some concrete cases. Consider, say, petitioning officials. Perhaps you care most about distributive equality. Officials make all kinds of decisions that affect such equality. You don't know anything about many of these decisions. But if you have an LLM representative that contacts an official and expresses a view in favor of the more egalitarian option in each specific case, that increases your impact on the level of distributive equality. You don't determine exactly what gets written in the letter, but your values contribute more to your ultimate end. Or consider elections. Suppose you care most about increasing economic growth. You might know nothing about the candidates running in local city or council elections. You might know very little about the candidates in European Union elections or for U.S. judicial posts. Thus, voting in these elections yields no knowing contribution to your ultimate end. But imagine your LLM representative carefully assesses your core values and then votes for the candidates who will promote them. Then these values have, through influencing the agent, influenced the character of your society.

This point also applies in direct democracies. Many polities, like California and Switzerland, routinely let voters vote directly on policies. This can increase their opportunity for authoring their political affairs. But most voters don't know much at all about the policies they are faced with in the ballot booth. They have little idea of the effects of different policies. In this context, giving people the ability to vote on such policies doesn't give them much more ability to influence things they care more about. If a voter really cares about, say, making housing cheaply available, but has no idea how rent control affects the housing market, then giving them the ability to vote on rent control does relatively little to contribute to their overall autonomy. But if a well-informed LLM representative carefully evaluates how a ballot measure will affect the achievement of what the voter cares most deeply about, then the voter, through controlling their LLM representative, can impact what their society looks like.

Let us sum up. Manuali may be correct that, if we create and empower LLM representatives, we enjoy little authorship of the exact issues the representative pushes for when we have not considered them. Our intentions are not linked to outcomes on these issues in the right way. But I suspect that voters see such exact issues as mere means to much broader, often vaguely defined, ends. They care less about what specific policies get enacted or about who in particular wins obscure elections than they do about the general features of their society. The central practical problem for ordinary voters is that they do not know how to use their vote to affect the general features of their society. LLM representatives can help solve this problem. Empowering them can let voters trade influence over the choice of means for that over choice of ends. Overall, this trade typically augments ordinary people's autonomy. If so, Manuali's pessimism comes from overlooking

the fact that we can exchange authorship over some things for authorship over other, more important, ones. My relative optimism is rooted in the idea that LLM representatives instantiate winning means-end trades for ordinary citizens.

There are practical upshots to this way of justifying LLM representatives. Existing prototypes of such representatives predict what people would say they prefer were they briefly asked about an issue (Gudiño-Rosero et al. 2024). Complementarily, some have suggested that the primary purpose of LLM representatives should be to allow policymakers to predict such expressed preferences (Lim and Savulescu 2026). This seems wrong to me. Instead, we should want LLM representatives to establish the values underpinning people's preferences and work out what policies best promote these values (cf. Lim and Savulescu 2026, 140–41). More specifically, we should want LLM representatives to work out what the citizen they represent cares most about, establish which political actions best contribute to the ultimate ends of their representee, and then take these actions. LLM representatives should not be tools of prediction for policymakers, but rather tools of influence for ordinary people.

4. Conclusion

None of this means that we should, all things considered, create and empower such LLM representatives. In my view, the most serious problem with doing so concerns equality: today, such representatives would very likely be based on models created by a small number of commercial companies. Empowering LLM representatives would thus give those who run these companies enormous political power. That would be objectionable by the lights of any broadly egalitarian justification for democracy (e.g. Christiano 2008; Kolodny 2014; Viehoff 2019). And there are technical problems concerning the accuracy of existing LLM-based agents. It is not obvious that existing LLM-based agents can effectively identify people's core values or the best means to achieving those values; if they cannot, they will not be very good tools of influence. Still, at the present moment, it seems to me unwise to bet against capability improvements in AI. My point, in any case, is narrow and theoretical: in principle, creating and empowering LLM representatives might augment rather than detract from citizens' autonomy.

Bibliography

Christiano, Thomas. 2008. *The Constitution of Equality: Democratic Authority and Its Limits*. Oxford University Press.

Converse, Philip. 2000. "Assessing the Capacity of Mass Electorates." *Annual Review of Political Science* 3: 331–53.

Gallo, Nick, and David E. Lewis. 2012. "The Consequences of Presidential Patronage for Federal Agency Performance." *Journal of Public Administration Research and Theory* 22 (2): 219–43. https://doi.org/10.1093/jopart/mur010.

Gudiño-Rosero, Jairo, Umberto Grandi, and César A. Hidalgo. 2024. *Large Language Models (LLMs) as Agents for Augmented Democracy*. July 30. https://www.alphaxiv.org/abs/2405.03452.

Han, Sung Jun. 2025. "The Question of AI and Democracy: Four Categories of AI Governance." *Philosophy & Technology* 38 (2): 83. https://doi.org/10.1007/s13347-025-00904-6.

Jarrett, Daniel, Miruna Pîslar, Michiel A. Bakker, et al. 2025. "Language Agents as Digital Representatives in Collective Decision-Making." arXiv:2502.09369. Preprint, arXiv, February 13. https://doi.org/10.48550/arXiv.2502.09369.

Kolodny, Niko. 2014. "Rule Over None II: Social Equality and the Justification of Democracy." *Philosophy & Public Affairs* 42 (4): 287–336.

Konya, Andrew, Lisa Schirch, Colin Irwin, and Aviv Ovadya. 2023. "Democratic Policy Development Using Collective Dialogues and AI." arXiv:2311.02242. Preprint, arXiv, November 3. https://doi.org/10.48550/arXiv.2311.02242.

Lavi, Liat. 2025. "One Person, One Bot." arXiv:2504.01039. Preprint, arXiv, March 31. https://doi.org/10.48550/arXiv.2504.01039.

Lim, James Edgar, and Julian Savulescu. 2026. "AI Preference Prediction and Policy Making." *AI & SOCIETY* 41 (1): 135–49. https://doi.org/10.1007/s00146-025-02474-9.

Lovett, Adam. 2020. "Democratic Autonomy and the Shortcomings of Citizens." *Journal of Moral Philosophy* 18 (4): 363–86.

Lovett, Adam. 2026a. "Is Decentralization Democratic?" *Critical Review of International Social and Political Philosophy*: 1-25.

Lovett, Adam. 2026b. "The Choice Argument for Proportional Representation." *American Journal of Political Science* 70 (3): 912–24. https://doi.org/10.1111/ajps.12980.

Lovett, Adam, and Jake Zuehl. 2022. "The Possibility of Democratic Autonomy." *Philosophy & Public Affairs* 50 (4): 467–98.

Manuali, Lorenzo. 2026. "When LLMs Threaten Democratic Autonomy." *Philosophy & Technology* 39 (3): 173. https://doi.org/10.1007/s13347-026-01184-4.

Viehoff, Daniel. 2019. "Power and Equality." In *Oxford Studies in Political Philosophy Volume 5*, edited by David Sobel, Peter Vallentyne, and Steven Wall. Oxford University Press.

Xu, Guo. 2018. "The Costs of Patronage: Evidence from the British Empire." *American Economic Review* 108 (11): 3170–98. https://doi.org/10.1257/aer.20171339.