

I hugely enjoyed reading this fascinating paper. I am thinking about how to represent policy decision-making in a global agent-based model myself and this work seems as though it could be the genesis of a substantial breakthrough in this area. My questions therefore primarily relate to scaling and generalisability. Please take them in a constructive and positive spirit – from someone who is trying to understand your work to apply it themselves.

Scaling

You note scaling issues in your discussion, however I think some further information and details would be helpful for the reader to gauge the extent of these challenges.

- It would be helpful to get a sense of run-time of one human-in-the-loop prompt session per institutional agent type (setup, LLM thinking time etc). E.g. Suppose we wanted to do sets of runs with differing levels of policy targets to assess consistency of answers, and /or to work towards an heuristic set of policy actions without human-in-the loop interventions. How feasible is this?
- Similarly, perhaps I missed this, but if you repeat one of your policy scenarios multiple times, how much do the resulting outputs diverge? Is this computationally prohibitive?
- Further, if you have stochasticity in the underlying model, how much can this lead to unpredictable policy pathways? You provided historical data to the LLM – I take it this was observations rather than historical CRAFTY runs? Did you then spin-up the CRAFTY AFT distribution to match the observations? Otherwise might our initial LLM choices be sensitive to the initial conditions?
- Overall, I think a supplementary information expanding lines 179-191 to be really clear on what you did would be very helpful and make the work more reproducible. E.g. which RCPs / SSPs did you use? I *think* from Figure 5 that you are showing means across the scenarios, but this isn't totally clear. It would also be very interesting to see how much variance there is by scenario.

Generalisability

You observe that stakeholder disagreement & subsequent contested policy spaces lead generally to slower decision-making. This is an important and fundamental insight, with some grounding in the literature. Some questions and comments below on how universal / generalisable such a finding may / may not be.

Lines 308-312

The setup of agent Q overall seems good and appropriate. I think it is worth being careful to remind readers that you are explicitly mimicking policymaking processes in a European context – with broadly democratic norms and systems. The text here seems to discuss multi-stakeholder policymaking in abstract terms, but the setup of the multi-stakeholder network would presumably have to vary substantially in other policymaking systems. For example, in

more authoritarian government systems, we may have “industrialisation from above” with very rapid changes, or one group of stakeholders’ rights and views being cut out of decision-making.

Lines 407-414

Here and elsewhere you state that slow and or incremental policy changes are more realistic / more in line with expectations than the optimisation algorithm. A few more references to support this would be good, particularly to clarify whether this is primarily a feature of western democratic systems or a more general phenomenon.

That said, let us assume that incrementalism is a broadly realistic simulated policymaking approach. I wonder if, in a subsequent paper, one could demonstrate this empirically?

E.g. could we take countries’ stated climate targets, and review concrete progress / policy implementation towards them vs what an economically-optimal trajectory towards achieving them might look like. If it could be clearly demonstrated that these simulated policy responses are closer to real-world choices than optimisation-based modelling that would be a tremendously important finding, I think. Not only to evaluate your model, but also for wider consideration of institutional constraints on rates of environmental land use change.