Towards psychological realism in agent-based models

Dr Toby Prike, University of Southampton

Hello, I'm Toby Prike and I'm going to be talking about increasing the psychological realism of agent-based models. In other words, making the agents within the models behave more like real humans. When building agent-based models in social science, it's important to keep in mind what the agents within the model represent. Most commonly, the agents are intended to represent humans and therefore, if the model is going to be a useful representation of the world, then the agents within it need to behave at least somewhat similarly to real humans.

Traditionally, many models have assumed that people behave rationally and optimally. For example, agents often have perfect and complete information, are unbiased when collecting information and updating their beliefs, and make optimal choices and decisions. However, there is considerable evidence from psychology and other social sciences that calls these assumptions into question. The assumptions made within a model will have major impacts on the results, so it is crucial that the assumptions made for agents are well justified and evidence based.

I will now run through a few examples of research findings that can potentially be applied to agent-based models to make them more psychologically realistic.

When making complex real-world decisions, people rarely have complete and accurate information. Therefore, it is important that this reality is represented within agent-based models. For example, in an agent-based model of migration, rather than assuming that all agents start off with perfect knowledge of potential migration destinations and routes, you can instead have agents start off with partial knowledge, and allow for some of this knowledge to be inaccurate or distorted. Agents can then gain knowledge about destinations and routes while the model runs, such as by exploring the model landscape or interacting with other agents.

This ties in well with another important area for improving the psychological realism of agent-based models, the process of learning new information and integrating this new information to update beliefs. Learning and updating beliefs are not simple or straightforward processes. People do not learn new information perfectly, with much of the information encountered quickly forgotten and never entering more durable, long term memory. Even new information that is initially learned and remembered will decline quite rapidly unless active steps are taken to relearn and consolidate the information in memory. Therefore, any model in which agents immediately and durably remember all information that they encounter will be modelling agents that have memories unlike that of human beings. And, assuming that the amount and accuracy of the information that agents possess influences the decisions and actions they take, this could potentially lead to considerable divergence between the agents in the model and the actual humans they are intended to represent. As well as having imperfect memory, people do not update their knowledge and beliefs in an optimal and unbiased way. For example, as anyone who has ever been in an argument or debate can tell you, people are usually highly resistant to changing strongly held beliefs, even when presented with considerable counterevidence. Therefore, agent-based models should take this into account and allow for factors such as error and bias when agents are gaining new information.

In an agent-based model of migration developed within our team, we have made attempts to add some psychological realism to the models in terms of knowledge and belief updating. Rather than starting with large amounts of information, within the model each agent starts off with some limited information and can then gain additional information as the model runs and they explore the landscape and communicate with other agents. Additionally, when gaining this new information, they do not learn it perfectly. They do not learn all information that they encounter and when they do learn information from other agents, random error is added to reflect the imperfect nature of learning and remembering information. Agents can also be biased in their assessment of new information. For each piece of information that agents have they also have a level of trust in that information. When agents exchange information, their new belief will be influenced by both the information exchanged and the level of trust that they and the other agent have in that information. Therefore, agents will be less willing to update their belief if they have a high level of trust in their initial belief. The trust levels of the two agents can also interact, leading to potential outcomes such as one agent being convinced that information from the other agent is trustworthy, neither agent being convinced, or one or both of the agents becoming more confused and reducing their trust in the information.

I have also conducted several psychological experiments with my colleagues to try and better inform agent-based models of migration. Recently, we conducted a study examining how people weight migration information from different sources. Participants were given information stating that a migration journey was very likely, likely, unlikely, or very unlikely to be safe. They received these pieces of information from five different sources, with one piece of information per source. The five sources chosen were, a news organisation, an official organisation, a family member, the person organising the travel, and someone with relevant personal experience. Participants gave judgements of the likelihood that the journey would be safe and made a yes/no decision of whether they would travel. We found that people gave more extreme responses when presented with information from an official organisation. That is, compared to other sources, when an official organisation said the migration journey is likely to be safe, participants gave higher likelihood of safety ratings, and when an official organisation said the migration journey is unlikely to be safe, they gave lower likelihood of safety ratings. We also found that official organisations had the most influence on people's overall likelihood of safety ratings and decisions to travel, judgements and decisions made after people had received all the pieces of information. The second most influential source for the overall ratings and decisions was people with relevant personal experience. These findings can be incorporated into an agent-based model by allowing some sources of information to more strongly influence agents' beliefs and decisions. Specifically, these results suggest official organisations and people with personal experience will have a stronger influence.

In another study we conducted, people chose between two potential migration destinations that varied in terms of potential gains and losses in monthly income. This allowed us to test whether, as is often found in financial domains, participants deviated from economic rationality. We found that, consistent with many other research findings, participants were loss averse, meaning they weighted the potential of losing income more strongly than potential gains in income of the same size. This finding differs from the rationality assumptions often used in models but can be incorporated into an agent-based model with relative ease. By allowing agents to weight losses more heavily than gains, the assumptions of the model will be more consistent with the research literature and more psychologically realistic.

Additionally, in this same study we found that participants showed decreased sensitivity for gains in income as they got further from the reference point. That is, as increases in monthly income got

further from their current income level, they placed less and less value on these increases in monthly income. This finding could again be incorporated into an agent-based model by using a function to decrease the weight of potential gains as they get further from either the agents' current situation or their current expected outcome, whichever is the more defensible choice of reference point.

As outlined thus far, it is vital that agent-based models of social systems strive to be more psychologically realistic. However, there are many challenges when it comes to adding psychological realism to agent-based models. One key challenge is that much of the social sciences literature is based on finding so-called statistically significant results that differ from the null. For example, showing that gender, treatment condition, or a personality trait has a non-zero effect on an outcome measure of interest. However, the range of plausible effect sizes for many published results often stretches all the way from near zero to quite large and attempts to accurately and precisely estimate effect sizes are rare. This means that even if there is a solid evidence base, for example that people are biased when updating their beliefs or that they weight losses more heavily than gains, it can still be difficult to apply these findings to agent-based and other models because of the need for formalism and setting values (or ranges or distributions of values) for inputs. Therefore, it is necessary to know not just that people are biased or loss averse, but also to know how biased they are and how loss averse. Do they weight losses twice as much as gains? One and a half times? Three times? What should be done if the estimates from different studies or methodologies are inconsistent?

There are also some more general and fundamental issues with applying results from experimental research to agent-based and other models. One issue that is relevant for nearly all psychological experiments, is how to simulate complex real-world environments within the confines of an online or lab-based experiment. For example, migration decisions have the potential to massively change one's life, involve incredible upheaval, and, in the case of forced migration, may occur in life-threatening circumstances. Any lab-based or online experiment cannot come close to replicating the real-world stakes or magnitude of these decisions and this is a major challenge for both designing experiments and for applying existing insights from the literature to agent-based models of real-world phenomena.

Another major challenge is that decisions are highly context dependent and influenced by a huge number of factors. Therefore, even if it is possible to gain insight into specific aspects of decision making, there remains the challenge of establishing the extent to which these insights are applicable in other contexts as well as the additional challenge of understanding and/or making reasonable assumptions about how these various factors interact. The extent of this challenge is made explicit when you are required to assign parameter values to multiple factors and formalise the interactions between these different factors within an agent-based model.

Although these are incredibly difficult challenges to overcome, they must be addressed, not only to better enable modelling in social science but because they are challenges that must be addressed for the sake of social science more broadly. If we can make progress in addressing these challenges, then we will be able to add much greater psychological realism to agent-based models. These developments will allow us to take greater advantage of the incredible potential that agent-based models have for improving our understanding of a vast number of social challenges and research questions.