Report on model development

Introduction

Climate lobbying, broadly defined as efforts to influence policies and public attitudes on climate change, plays a crucial role in shaping both environmental outcomes and public discourse (Yeganeh, 2020; Schlichting, 2013). Over decades, climate lobbying has evolved from grassroots advocacy to a multifaceted arena of influence, encompassing diverse actors such as environmental NGOs, corporate interests, and international coalitions (Dunlap, 2013; Robinson, 2006; Stokes, 2017). These efforts have profound implications for the adoption of climate policies, the regulation of emissions, and the promotion of renewable energy (Nisbet, 2014; Rowell, 2013).

Despite its prominence, the integration of lobbying dynamics into formalized opinion dynamics models remains an underexplored area. Opinion dynamics models, widely utilized across disciplines, provide a robust framework for understanding how beliefs evolve within heterogeneous populations (Hegselmann and Krause, 2002; Sobkowicz, 2020). These models, which draw on concepts from social influence (Centola et al., 2007), statistical physics, network theory, and behavioral economics, offer valuable insights into the mechanisms of consensus formation and information diffusion (Das et al., 2014; Li et al., 2020). Their application to climate lobbying presents a promising avenue to analyze the interplay between lobbying strategies, public perception, and policy decisions (Sirbu et al., 2013; Sirbu et al., 2017).

The **ALMONDO project** bridges these gaps by developing and analyzing an opinion dynamics model that incorporates behavioral biases and strategic lobbying. Its second activity is specifically dedicated to developing a theoretical model that formalizes biases in individual learning processes, applied to climate change. This approach provides fresh perspectives on the interactions between public opinion, behavioral learning, and climate lobbying. The modeling effort consists of three key tasks that collectively build an innovative framework to study how opinion dynamics evolve under the influence of behavioral biases and strategic lobbying activities.

The first task entailed conducting a comprehensive literature review to identify the behavioral biases most relevant to climate-related interactions. Here we report the findings of that review. Building on the insights provided by Botzen et al. (2021), our investigation highlighted the significance of directional and credibility-motivated reasoning (Druckman and McGrath, 2019). We incorporated these insights into a learning framework that accounts for underreaction to signals (Epstein et al., 2010; Massari, 2020; Bottazzi et al., 2023).

The second task developed a mathematical model of opinion dynamics, wherein agents with these behavioral biases interact within a network and are influenced by lobbyists. The framework is inspired by the Continuous Opinions and Discrete Actions model (Martins, 2008) and incorporates behavioral learning dynamics. Agents evaluate the likelihood of climate-related events based on probabilistic models, updating their beliefs upon receiving signals. Their updates are skewed by behavioral biases. These biases shape their responses to new information, tilting updates along the dimensions of directional reasoning and the credibility of sources.

Lobbyists, modeled as external agents who do not change opinion, aim to shape public beliefs by strategically disseminating information in support of specific probabilistic models. Operating under budget constraints, they engage in costly signaling to minimize the average Kullback-Leibler divergence between individual beliefs and their preferred model. This dynamic highlights the strategic nature of lobbying efforts in swaying public opinion.

The third task focused on analyzing the model's properties, both theoretically and computationally. Analytical exercises investigated two scenarios: one with a generic number of agents and a single lobbyist interacting over multiple periods, and another with two agents and two competing lobbyists interacting over one period. While the first exercise highlights the basic incentive structure and constraints of lobbying activities, the second one underscores their strategic interdependence and the broader implications for belief formation and consensus-building. Several numerical exercises were performed to assess the properties of the model in more complex scenarios. We simulated cases with zero, one and two opposing lobbyists. In each case, we considered two population structures: one with homogeneous agents in terms of under-reaction and no motivated reasoning, and another with non homogeneous agents and motivated reasoning.

Model characteristics

In this section, we describe the model introduced, concentrating on three important elements: (1) behavioral features that characterize the agents populating our models (connecting them with the evidence in the literature that drove our modeling choices), (2) agent interaction, and (3) lobbyists behavior.

Behavioral Learning on climate change

Our exploration of the literature identified five key contributions that shaped the way in which individuals in our model behave. A general perspective is provided by Botzen et al. (2021). Indeed, the authors individuate and discuss six behavioral biases that affect individuals when dealing with low-probability high-consequences risks, such as climate change. They are *simplification*, *availability*, *a finite pool of worries*, *myopia*, *"not in my term of office" bias*, and *herding*. *Simplification* refers to the tendency to concentrate on either the chance of a climate-related disaster happening or the severity of its possible outcomes, instead of making a rational choice involving both dimensions. *Availability* concerns agents underestimating the risk of a climate-related event until they are directly or indirectly affected (e.g., through social connections). *Finite pool of worries* refers to individuals having a limited amount of "emotional resources" to be allocated to sources of concern, hence, when urgent worries are experienced, climate change concern is reduced. *Myopia* refers to the tendency of evaluating decisions (such as investments or economic policies for climate mitigation) on a time span that is too short. The *"not in my term of office" bias* refers to the tendency of politicians to delay necessary but unpopular decisions regarding low-probability negative events, fearing the adverse impact on electoral support. *Herding* refers to the large influence that social connections have on the behavior of individuals. Among these six biases, those that are relevant for our analysis are simplification, availability, and herding; since to capture myopia, the finite pool of worries, and the "not in my term of office" bias one should consider a more general model with different risks, policies, and politicians. In our model, simplification occurs as agents base their decisions to signal for or against climate change solely on their subjective probabilities. Availability and herding follow, instead, from the network structure. Indeed, we assume that agents are embedded in a social network and update their beliefs on the basis of the signals they receive from their acquaintances.

On top of the aforementioned three biases, we considered the discussion provided by Druckman and McGrath (2019) on the presence of *motivated reasoning* in the process of climate change belief formation. Motivated reasoning indicates that the way in which agents process information depends on its goals. In particular, the authors distinguish between *accuracy motivation* and *directional motivation*. Accuracy motivation means that agents strive to be as accurate as possible. For instance, they tend to have high confidence in the results of a scientific study if the prior belief about the reliability of the scientists involved is high. On the contrary, directional motivation indicates that agents want to get to a given predetermined conclusion, hence, they interpret information in an asymmetric way depending on whether it agrees or disagrees with prior beliefs. While acknowledging the vast evidence pointing to motivated reasoning in belief formation processes on climate change, the authors argue that empirical and experimental studies in the literature do not allow one to distinguish between accuracy and directional motivation because of a form of observational equivalence. For the sake of our analysis, we focused on directional motivation. Indeed, introducing accuracy motivation means equipping agents with a trustworthiness belief distribution over the agents connected to them, including lobbyists. This introduces a layer of complexity and discretionality to the analysis that we preferred to avoid.[1] Directional motivation, as well as being widely invoked in the climate change literature (see, e.g., Hart and Nisbet, 2012), can be easily included in our model assuming that prior beliefs influence the way in which those priors are updated in posteriors.

Finally, we considered three contributions on learning with under-reaction (Epstein et al. 2010; Massari 2020; Bottazzi et al., 2023). These contributions provided us with the basic mathematical framework needed to model the belief updating process with the aforementioned cognitive biases. The key idea is that agents build their (subjective) conditional probabilities by means of a convex combination of the conditional probabilities of some models (i.e., probability distributions over future events). The weights used to combine the models can be thought of as the probabilities assigned to each model to be the correct one. Those weights evolve over time depending on the sequence of signals observed and the way in which they are updated manifests the biases of the agents. In particular, under-reaction means that agents convexly combine the Bayesian update with the prior, such that they react to new information less than what a pure Bayesian agent would have done. The quantity determining the convex combination is generally referred to as λ, and we explored two scenarios for it. In the first scenario, it is an agent specific (constant) parameter. In the second scenario, it is a function of priors and the nature of the signal. Hence, if a signal favoring a given model is received and the agent assigns high confidence to such a model being the correct one, then it strongly updates weights in favor of the model. On the contrary, if the received signal favors a model to which the agent attaches very low confidence, then weights remain (almost) untouched. We further assumed that an agent specific parameter α gauges the strength of such a bias. These features directly capture the idea of directional motivated reasoning. The mathematical equations describing the evolutions of beliefs involve parameters that shall be calibrated during the model calibration activity exploiting the results of the experiment.

[1] Nonetheless, this paves the way to a future contribution in which we extend the model to consider accuracy motivation.

Agents’ Interaction and Lobbyists

Interaction among agents is mediated by a network of connections. We assumed that such a network is exogenously provided, as it will be shaped during the calibration phase of the project to respect the properties empirically observed during the data mining activity. Over such a network, agents send and receive signals for a finite number of periods. Signals can be of two types: either supporting the occurrence of a negative climate-related event after the finite number of interaction periods, or dismissing it. For the sake of our analysis, we assumed that the agents form their beliefs combining two probabilistic models, a “pessimist” (or “realist”) one and an “optimist” one. The key difference between the two models is that the first one attaches a larger probability to the negative event to occur than the second one. Hence, in each interaction period, one of the agents in the network is randomly selected to communicate to all of its connections in this round. The type of the signal (either supporting the occurrence of the event or dismissing it) is drawn according to the subjective conditional probability distribution of the agent. In this way, we capture the simplification bias described above assuming a sort of discrete-choice decision rule. The signal is sent to all the individuals connected to the communicating agent and their subjective probabilities are updated according to the behavioral learning procedure illustrated above. In the baseline version of the model, lobbyists are not active and, as the rounds of interaction go by, agents’ subjective probabilities evolve assimilating the signals received from connections. A numerical exploration of the evolution of opinions (i.e., subjective conditional probabilities) is provided at the beginning of the “Numerical Exercises” section.

Then, lobbyists are added to the model. They can be thought of as external agents with fixed opinions that, in every period, can send signals to the individuals communicating on the social network. Each lobbyist is assumed to support one of the two models. In each period, lobbyists send signals in favor of the supported model to individuals. Those signals are processed by individuals in the social network in the same manner of signals coming from peers. Hence, they update beliefs according to the procedure described above. The objective of each lobbyist is to minimize the expected average relative entropy (or Kullback-Leibler divergence) of the final individual beliefs with respect to the supported model. Sending a signal is costly and lobbyists must take their communication decisions under a budget constraint. For the sake of simplicity, we assumed that the cost of sending a signal is constant and uniform. Notice that, on the one hand, as the number of active lobbyists is larger than one, a strategic interaction framework naturally emerges. Indeed, we exploited notions and techniques from game theory to study the equilibrium behavior of lobbyists in a special case (see the section “Theoretical Exercises” below). On the other hand, in the “Numerical Exercises” section, we performed several simulations assuming that lobbyists randomly draw their communication strategy. In this way, we could better understand how the model reacts to different incarnations of the lobbying activity. Notice that, during the model calibration activity, we shall use the results of the data mining activity on lobbies in Europe to calibrate lobbying behavior in terms of communications strategies and budgets.

Theoretical Exercises

In this section, we propose two theoretical exercises in special cases that shed some lights on some basic properties of the model. In the first exercise, we analyzed a situation in which only one lobbyist is active in a rather generic social network. This unveils the basic incentive structure of the lobbyists and the effect of the budget constraint. In the second exercise, instead, we assumed that two lobbyists (supporting opposite models) are active in a social network populated by two individuals (an “influencer” and a “follower”) learning in a Bayesian way (neither under-reaction nor directional motivated learning are active) and interacting for one period. This unveils the effects of strategic interaction in the simplest scenario.

One Lobbyist

In this exercise, a generic number N of individuals interact on a generic social network for a number T of rounds. Agents are characterized by behavioral biases in learning according to the discussion above. At the same time, for this particular exercise, it is not important whether the parameter characterizing the convex combination of the prior with the Bayesian update in the updating rule is a constant parameter or displays the directional motivated reasoning. Indeed, the important feature is that, as a signal from the lobbyist is received by the agent, the conditional probability of the individual moves towards the model supported by the lobbyist. Hence, the constraint minimization problem of the lobbyist can be easily set up considering the expected average relative entropy of final beliefs with respect to the supported model as objective function and the budget constraint composed by the value of signals sent smaller or equal than the budget. Given the properties of the relative entropy, one immediately notices that the objective function is strictly decreasing in the number of signals sent. Hence, the optimal communication strategy of the lobbyist is always to try to exhaust the budget constraint sending signals to agents. More specifically, we need to distinguish between two cases. In the first one the budget is so large that the lobbyist can send a signal to each agent in each period. In this case the optimal strategy is to communicate to each agent in each period. In the second case, instead, the budget is not enough to always signal to everyone. In this scenario, every optimal strategy will involve sending a number of signals equal to the floor of the budget divided by the unitary cost of a signal. At the same time, deciding to which agents and in which periods to send signals is not straightforward, since the interaction between the behavioral biases and the communication among agents complicates the analysis. In the special case in which the network is composed by only one agent, then the timing of the signals becomes unimportant and every communication strategy that exhausts the budget constraint is optimal.

Two Agents and Two Lobbyists

In this exercise, on the one hand, we simplified the social network structure, while, on the other hand, we considered two lobbyists supporting opposite models. In this way, we were able to flesh out the strategic interaction that characterizes lobbying activity. The social network involves only two agents. One of them, identified as agent A, is an "influencer", while the other one, agent B, is the "follower". Indeed, the only ex-ante difference between the two is that A always communicates to B. Both agents are Bayesians, hence, no bias is active for any of them. This is a simplifying assumption that allows one to easily compute subjective conditional probabilities. At the same time, considering the framework without biases acts as a valuable benchmark for the rest of the analysis. Individuals interact for only one round and lobbyists have enough budget to send only one signal. Lobbyists can randomize their choices, hence, their problem is to choose the mixed strategy that is a best reply to the decision of the rival. A mixed strategy, in this setting, is a probability distribution over the two actions "send a signal to A" and "send a signal to B". The payoff of the ensuing game is simply the expected average relative entropy emerging from the combination of mixed strategies chosen by the lobbyists.

The outcomes of the game can be studied by means of the Nash equilibrium concept: a Nash equilibrium is a combination of mixed strategies (one for each lobbyist) such that no one has incentive to deviate (see Mas-Colell et al., 1995). This can be easily done by analyzing a plot showing, for both lobbyists, the best reply probability that one lobbyist attaches to send a signal to agent A as a function of the probability that the opponent attaches to the same action. The points in which the functions cross are the Nash equilibria of the game. Figure 1 shows such a plot for a situation in which the pessimist and the optimist models are extreme and opposite. In such a scenario, a unique Nash equilibrium emerges: both lobbyists send a signal to agent *A*, the influencer.

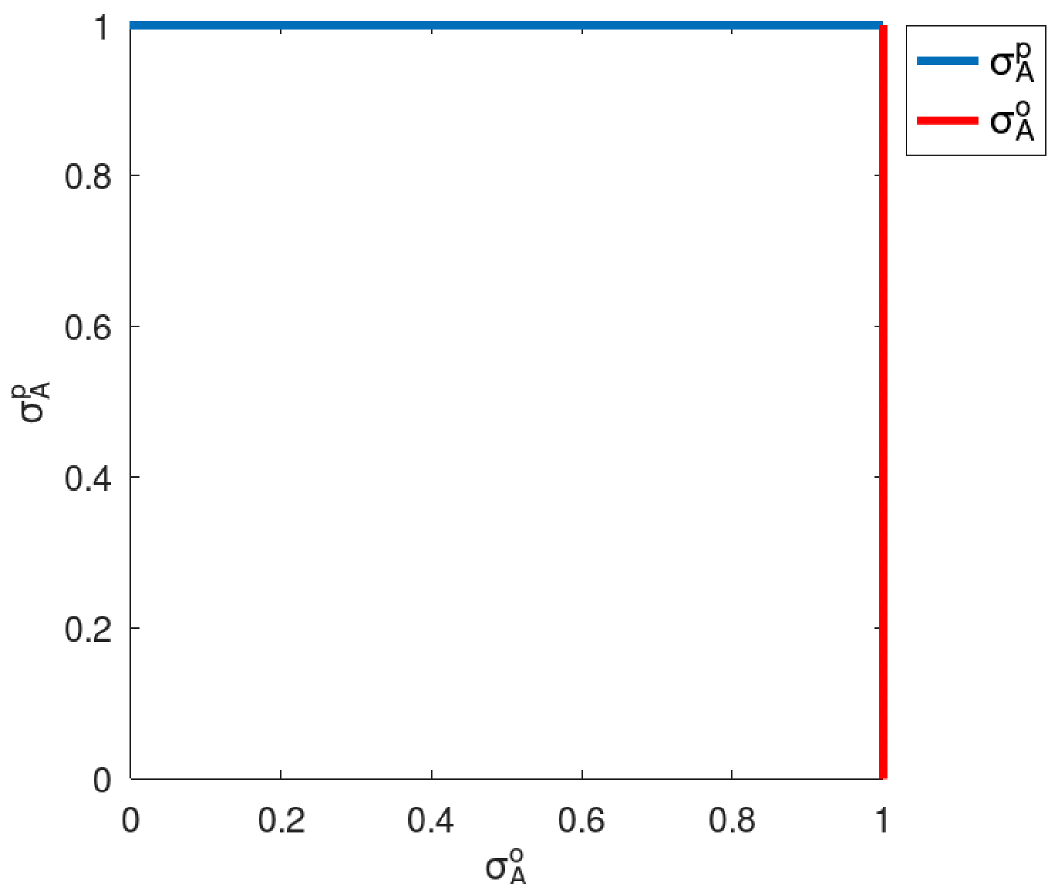

Figure 1 Mixed strategies of the lobbyists for the case in which the probability of the negative climate event is 0.99 for the pessimist model and 0.01 for the optimist model. σ_A^o is the probability the lobbyist supporting the

optimist model attaches to send a signal to A in its mixed strategy, σ_A^p is the probability the lobbyist supporting the pessimist model attaches to send a signal to A in its mixed strategy.

The same conclusion is obtained in the case in which the probability attached to the negative climate event is rather extreme for one of the two models and intermediate for the other. This clearly appears in Figure 2, where we show the best reply mixed strategy profile in the case in which the probability attached to the negative event is 0.5 for the pessimist model and 0.1 for the optimist model and in the case in which the probabilities are 0.9 for the pessimist model and 0.5 for the optimist model.

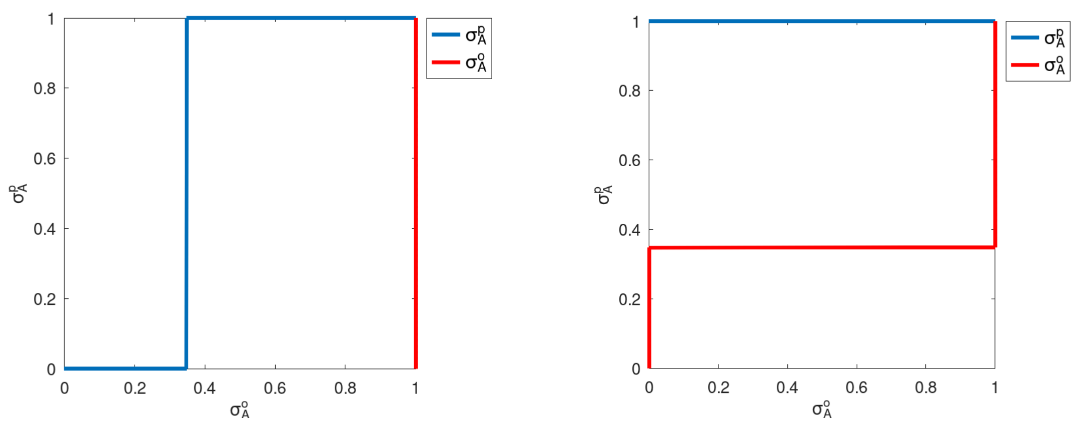

Figure 2 Mixed strategies of the lobbyists for the case in which the probability of the negative climate event is 0.5 for the pessimist model and 0.1 for the optimist model (left) and for the case in which the probability of the negative climate event is 0.9 for the pessimist model and 0.5 for the optimist model (right). σ_A^o is the probability the lobbyist supporting the optimist model attaches to send a signal to A in its mixed strategy, σ_A^p is the probability the lobbyist supporting the pessimist model attaches to send a signal to A in its mixed strategy.

The reason for the outcomes observed in Figures 1 and 2 is that, for at least one of the two lobbyists, sending a signal to the "influencer" is always the best choice.

When, instead, the probabilities attached to the event by the models tend to be close and asymmetric, the only Nash-equilibrium mixed-strategy profile involves both lobbyists assigning non-zero probabilities to both actions. Figure 3 supports such a point showing the best reply mixed strategy profile when the probability attached to the negative event is 0.99 for the pessimist model and 0.9 for the optimist model and when the probabilities are 0.1 for the pessimist model and 0.01 for the optimist model.

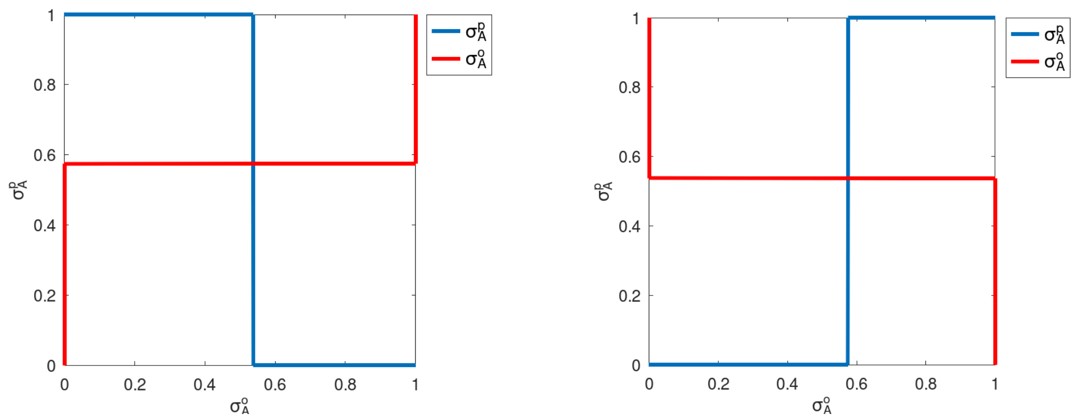

Figure 3 Mixed strategies of the lobbyists for the case in which the probability of the negative climate event is 0.99 for the pessimist model and 0.9 for the optimist model (left) and for the case in which the probability of the negative climate event is 0.1 for the pessimist model and 0.01 for the optimist model (right). σ_A^o is the probability the lobbyist supporting the optimist model attaches to send a signal to A in its mixed strategy, σ_A^p is the probability the lobbyist supporting the pessimist model attaches to send a signal to A in its mixed strategy.

Finally, in Figure 4, we present the strategic interaction emerging when the probabilities of the models are intermediate and close (probability of 0.55 of observing the negative event under the pessimist model and probability of 0.45 of observing the negative event under the optimist model). Here, three Nash equilibria emerge, in two of them the lobbyists coordinate in sending the signal to the same individual (either the "influencer" or the "follower"), in the remaining one they randomize their choices.

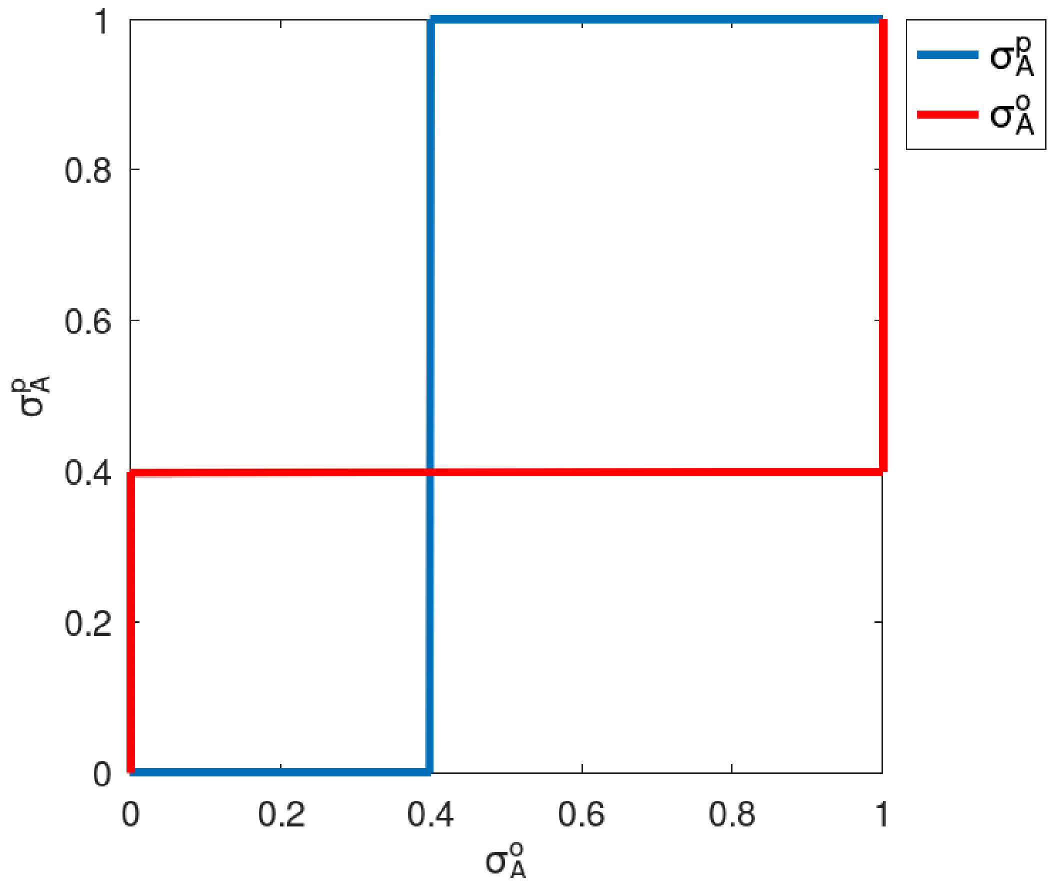

Figure 4 Mixed strategies of the lobbyists for the case in which the probability of the negative climate event is 0.55 for the pessimist model and 0.45 for the optimist model. σ_A^o is the probability the lobbyist supporting the optimist model attaches to send a signal to A in its mixed strategy, σ_A^p is the probability the lobbyist supporting the pessimist model attaches to send a signal to A in its mixed strategy.

Notice that in the cases of Figures 3 and 4, somehow counterintuitively, lobbyists may have incentive to signal to the "follower". This occurs because, depending on the probabilistic environment they face, they may have incentive to either copy the behavior of the opponent or to differentiate from it. This is particularly clear in Figure 3. Indeed, we have that for one lobbyist it is optimal to assign zero probability to signalling to the "influencer" when the probability attached to the same strategy by the opponent is low, while it is optimal to attach full probability to it when the opponent assigns it a sufficiently high probability. For the other, instead, it is the other way round. Interestingly enough, the lobbyist that would like to differentiate is the one supporting the pessimist model when models' probabilities are close to 1, while it is the other one when probabilities are close to 0. Hence, we can infer that two forces are at play: one working to offset the signaling action of the rival influencing the same agent, and the other striving to differentiate in order to ensure an effect on the opinions of an individual. The interplay of these forces may create incentives for one (or both) lobbyists to signal to the "follower".

Following the discussion of Botzen et al. (2021), climate-related negative events are *low-probability–high consequence* events. Hence, among the cases presented, the one shown in the right panel of Figure 3 appears the most realistic one. This indicates that we expect "rational" lobbyists to randomly choose between signalling to the "influencer" and the "follower" according to a probability distribution that is slightly skewed towards sending a signal to the "influencer".

Numerical Exercises

In this section, we discuss a set of numerical simulations in order to describe the basic dynamics of the model. Results are displayed in the subsections below using different criteria: we are interested in the evolution in time of subjective opinions, and in the distribution of opinions during the simulation and after reaching the steady state. In all cases and scenarios, at the beginning of the simulation the agents' subjective probabilities are uniformly distributed in [0, 1] (see Figure 5).

First, we focus on simulations without lobbyists, aiming to showcase the model's benchmark behaviour before introducing lobbyist strategies. We consider two cases: the case with homogeneous agents (i.e the degree of underreaction λ is equal for all the agents) and no motivated reasoning, and the case with heterogeneous agents (i.e the degree of underreaction λ can be different between agents) and motivated reasoning.

The two cases display different asymptotic behaviour. In particular, in the case of homogeneous agents, for a number of iterations sufficiently large, subjective probability distributions collapse to a single realisation (see Figures 6 and 9). The speed at which the convergence takes place depends on the degree of underreaction (λ) set exogenously for all the agents: the higher the degree of underreaction (λ closer to 1), the longer the convergence time to a steady state.

Besides affecting the speed of convergence, the degree of underreaction has an effect on the dynamics of the transient period. In fact, when the agent is closer to bayesian rationality (λ closer to 0), the subjective probability distribution tends to polarize during the transient phase, as agents who receive a given signal suddenly and strongly update their own beliefs in the direction of the new information (Figure 7). After some iterations of "cycling" through the competing models, beliefs quickly converge to a consensus. In contrast, when agents have a higher degree of under-reaction (λ closer to 1) the transition to the steady state is slower and smoother (Figures 9 and 10).

In the case with motivational reasoning, where the underreaction parameter λ is conditional on the characteristics of the individual agent and of the signal received, the subjective probability distribution converges quickly to a Bernoulli distribution, as agent's beliefs polarize towards the optimist or the pessimist model (Figures 12-14). This is a very interesting behavior, as it is a more realistic case: rarely a population will converge on a single opinion, but generally clusters form. Therefore we have strong indications that the motivated reasoning mechanism we implemented is key to more realistic simulations.

Next, we introduce one lobbyist in the picture. At this stage, the "strategy" adopted by lobbyists is uniformly random, although in the future we aim to introduce more sophisticated and properly strategic behaviour into the picture. More in detail, the lobbyist decides to send a signal or not to send a signal with a probability of 50% for each agent. This means that for each time step, on average, it will send a signal to 50% of the agents shifting their beliefs towards their supported model. We also assume that the budget of the lobbyist is enough to support such a communication strategy. If the number of steps is T and the number of agents is N, with T sufficiently large, the budget of the lobbyist must be (approximately) TN/2.

The introduction of a single lobbyist in the model at this stage seems to drastically speed up the convergence of the simulation towards the outcome supported by the lobbyist. This comes as no surprise, as the introduction of lobbyists with a randomly generated strategy amounts to uniformly increasing the number of signals received on average by the agents in support for one model with respect to the other (Figures 15 and 16). Note that with the current configuration, the network emits one signal per timestep, while a single lobbyist emits N/2 signals and therefore is able to move the distribution towards their supported model very quickly. Interestingly, even when the motivated reasoning bias is considered, the signals sent by the lobbyist are strong enough to stir the network toward their preferred model and achieve a consensus on it (Figures 17 and 18).

Finally, we take into consideration the case with two lobbyists, both with the same uniform strategy, but supporting opposing models (Figures 19-22). In this case we see that the impact of the two lobbyists balances out, as the simulations display the same asymptotic behaviour as in the case with no lobbyist, both with homogeneous and heterogeneous agents. The only noticeable difference is in the speed of convergence, as simulations take less iterations to converge. This result should come as no surprise, as introducing lobbyists in the picture increases the average number of signals emitted per time step from 1 to T + 1.

Baseline Scenario, homogeneous agents, λ= 0.2

In the following numerical exercise, the probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model. Underreaction is low.

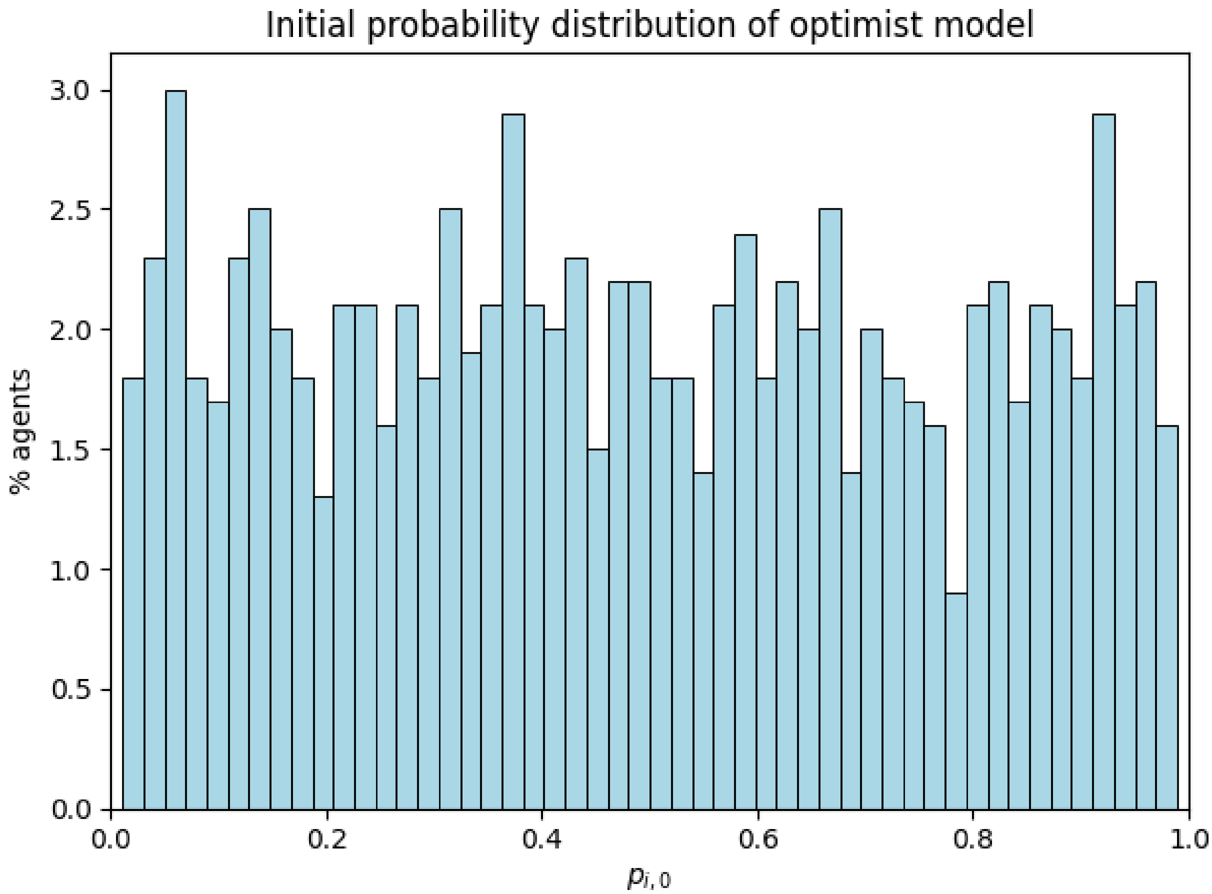

Figure 5 Subjective probability distribution at t = 0 for a typical run of the model. Initially, the agents' subjective probabilities are uniformly distributed. This initial distribution is common to both cases (both homogeneous and heterogeneous agents) and all scenarios inside them.

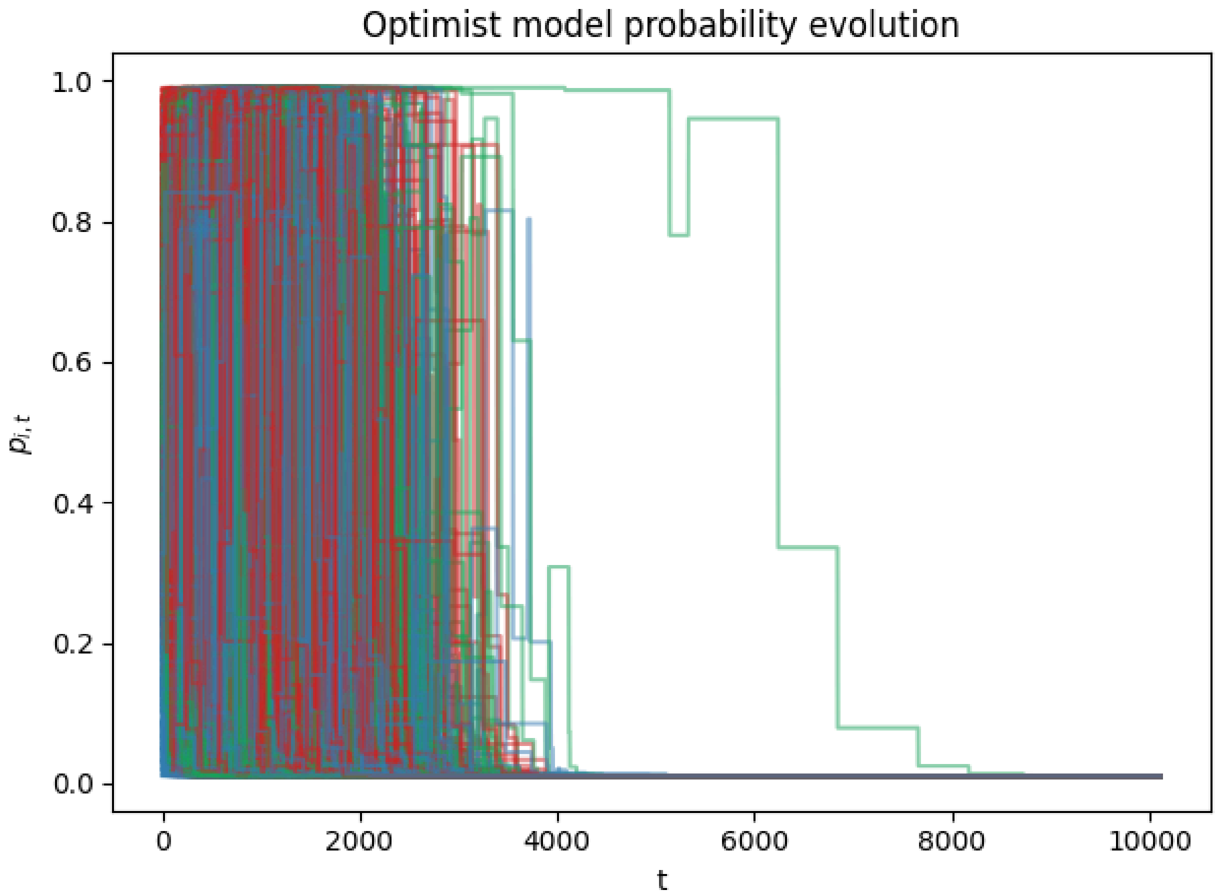

Figure 6: Evolution plot of agents' subjective probabilities for a typical run of the model with homogeneous agents (underreaction parameter $\lambda = 0.2$). The probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model. Each line in the plot corresponds to one agent, who changes their subjective probability while interacting with others. The colour encodes the initial agent opinion: from blue to red for subjective probabilities from 0 to 1. We note that the opinions converge to one value only.

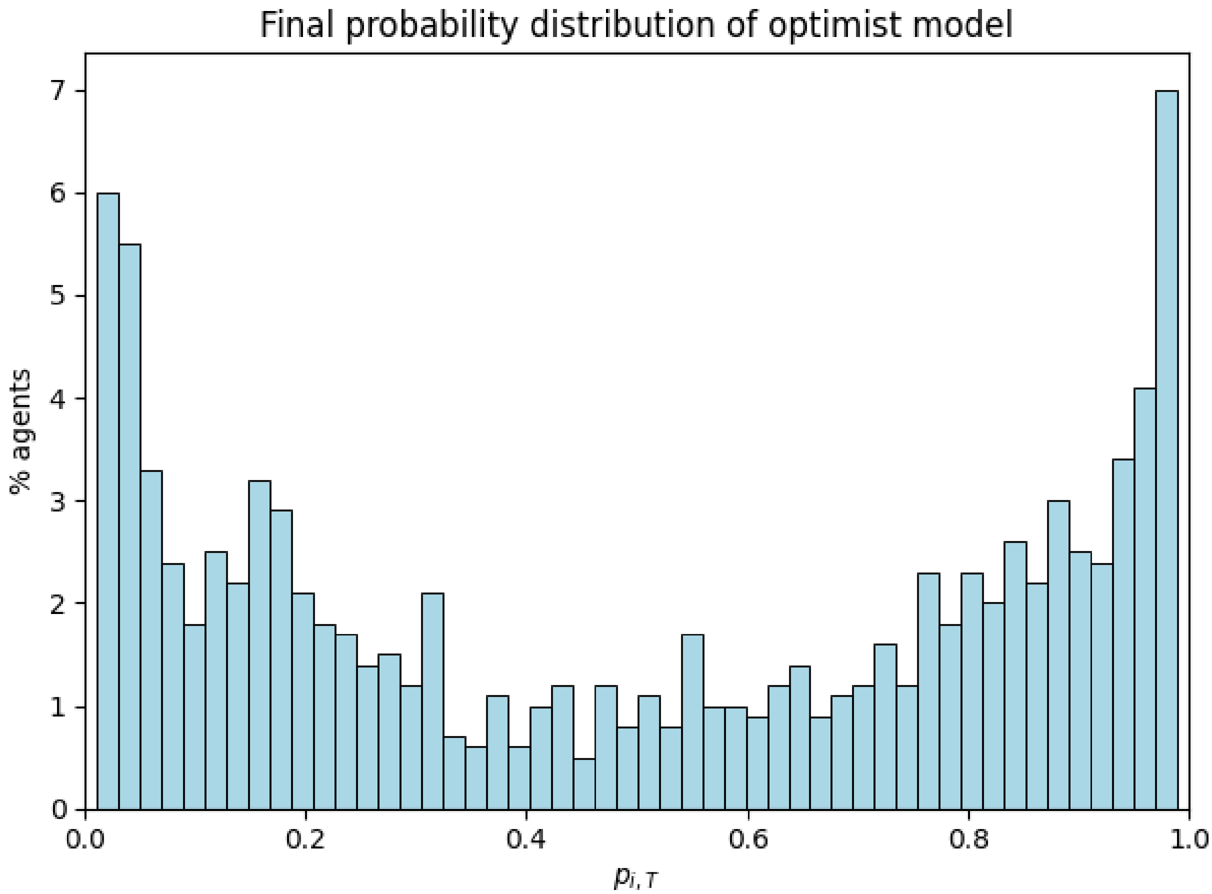

Figure 7 Subjective probability distribution at round t = 100 for a typical run of the model with homogeneous agents (underreaction parameter λ = 0.2). The probability of the negative climate event is p_p = 0.99 for the pessimist model and p_o = 0.01 for the optimist model.

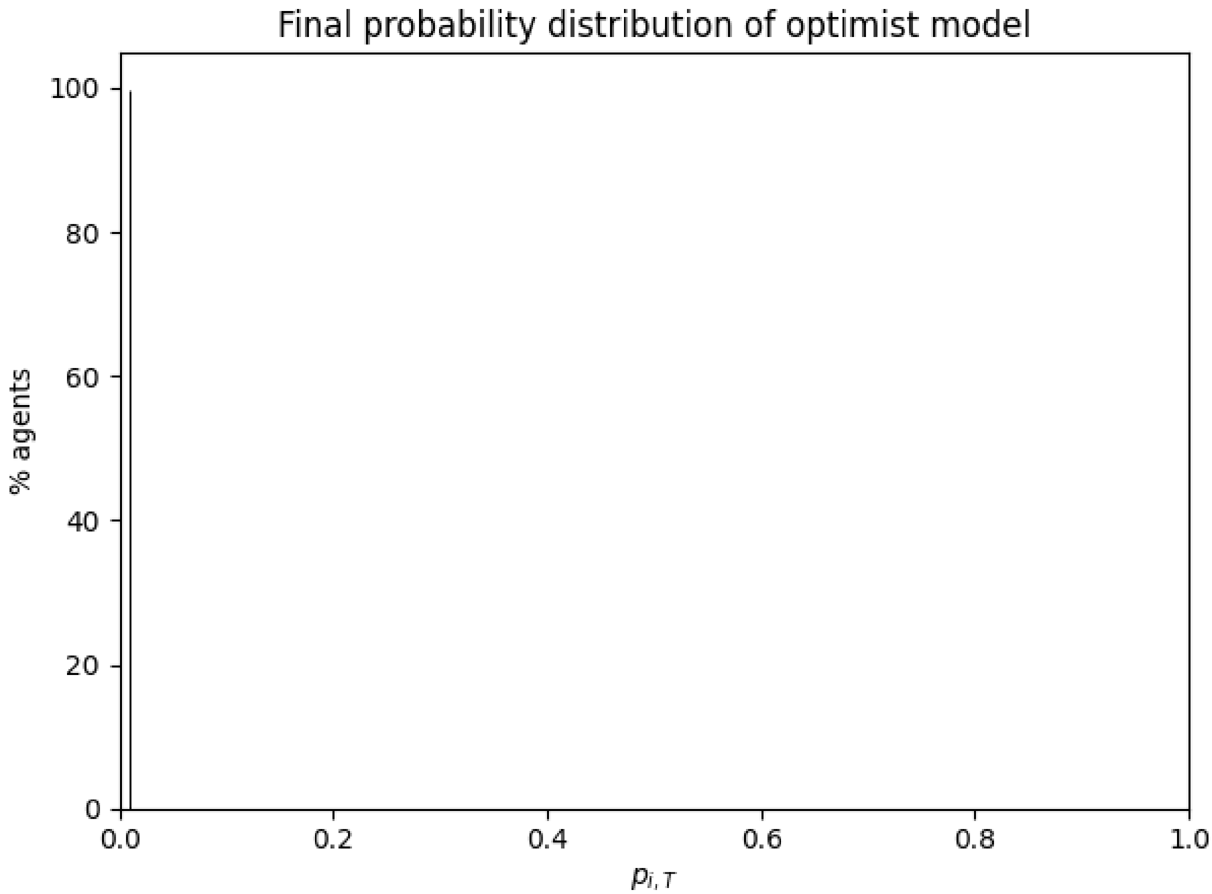

Figure 8 Subjective probability distribution at the steady state for a typical run of the model with homogeneous agents underreaction parameter λ = 0.2). The probability of the negative climate event is p_p = 0.99 for the pessimist model and p_o = 0.01 for the optimist model.

Baseline Scenario, homogeneous agents, λ = 0.8

In the following numerical exercise, the probability of the negative climate event is p_p = 0.99 for the pessimist model and p_o = 0.01 for the optimist model, similar to the previous case. However we increase the underreaction parameter.

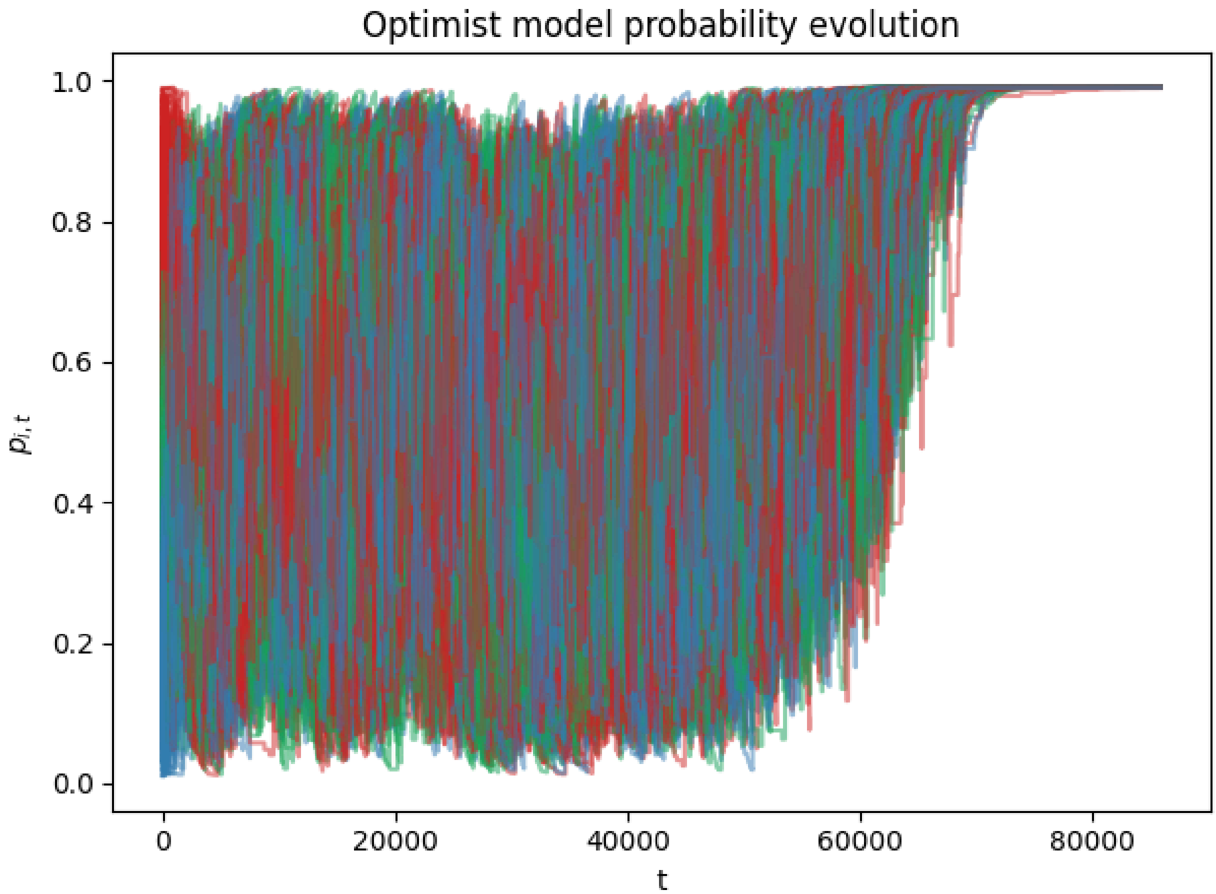

Figure 9 Evolution plot of the agents' subjective probabilities for a typical run of the model with homogeneous agents (the degree of underreaction is $\lambda = 0.8$ for all agents over time). The probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model;

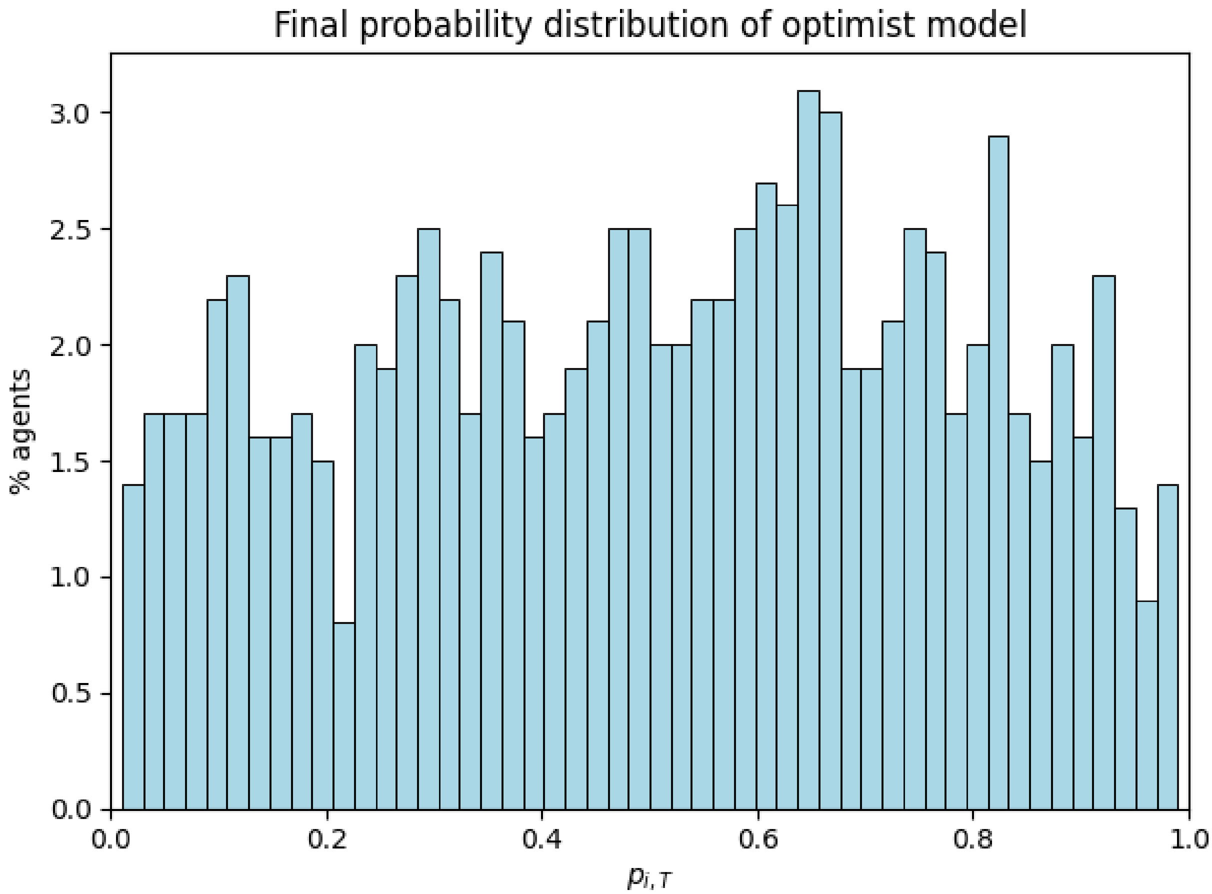

Figure 10 Subjective probability distribution at round t = 100 for a typical run of the model with homogeneous agents (underreaction parameter λ = 0.8). The probability of the negative climate event is p_p = 0.99 for the pessimist model and p_o = 0.01 for the optimist model.

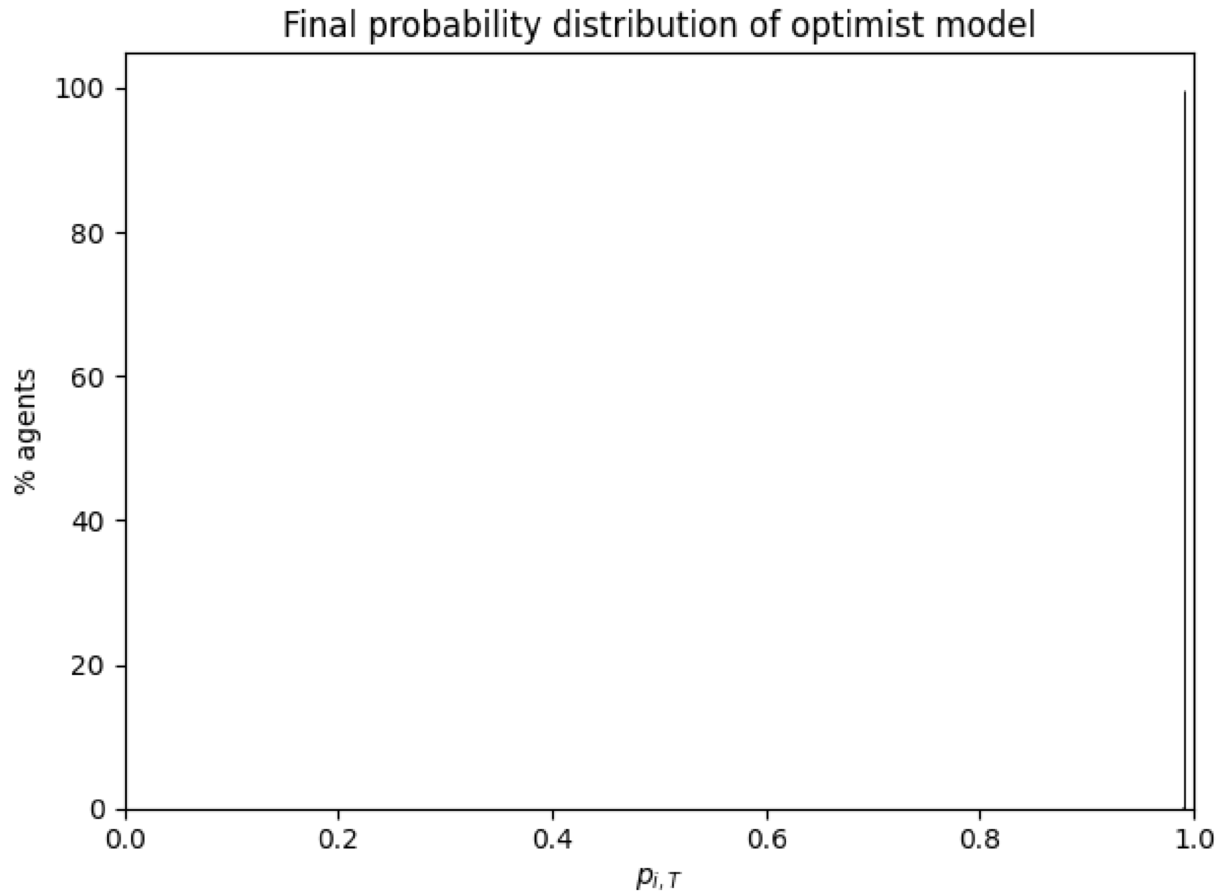

Figure 11 Subjective probability distribution at the steady state for a typical run of the model with homogeneous agents (underreaction parameter $\lambda = 0.8$). The probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model.

Baseline Scenario, heterogeneous agents, $\alpha = 0.5$

In the following numerical exercise, the probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model, however the λ parameter is not fixed but calculated for each agent using motivated reasoning, controlled by a new parameter α, controlling the strength of directional reasoning.

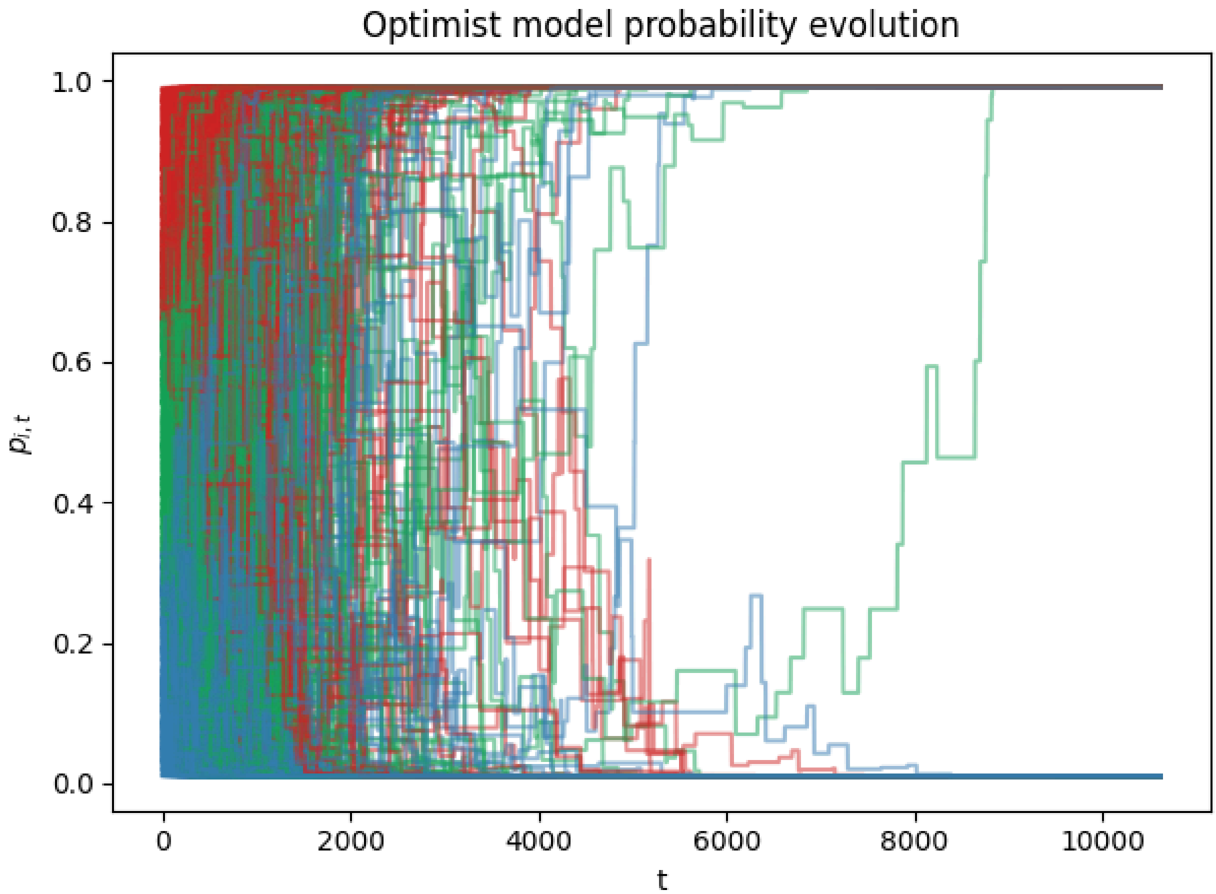

Figure 12 Evolution plot of agents' subjective probabilities for a typical run of the model with heterogeneous agents (directional reasoning parameter $\alpha = 0.5$). The probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model.

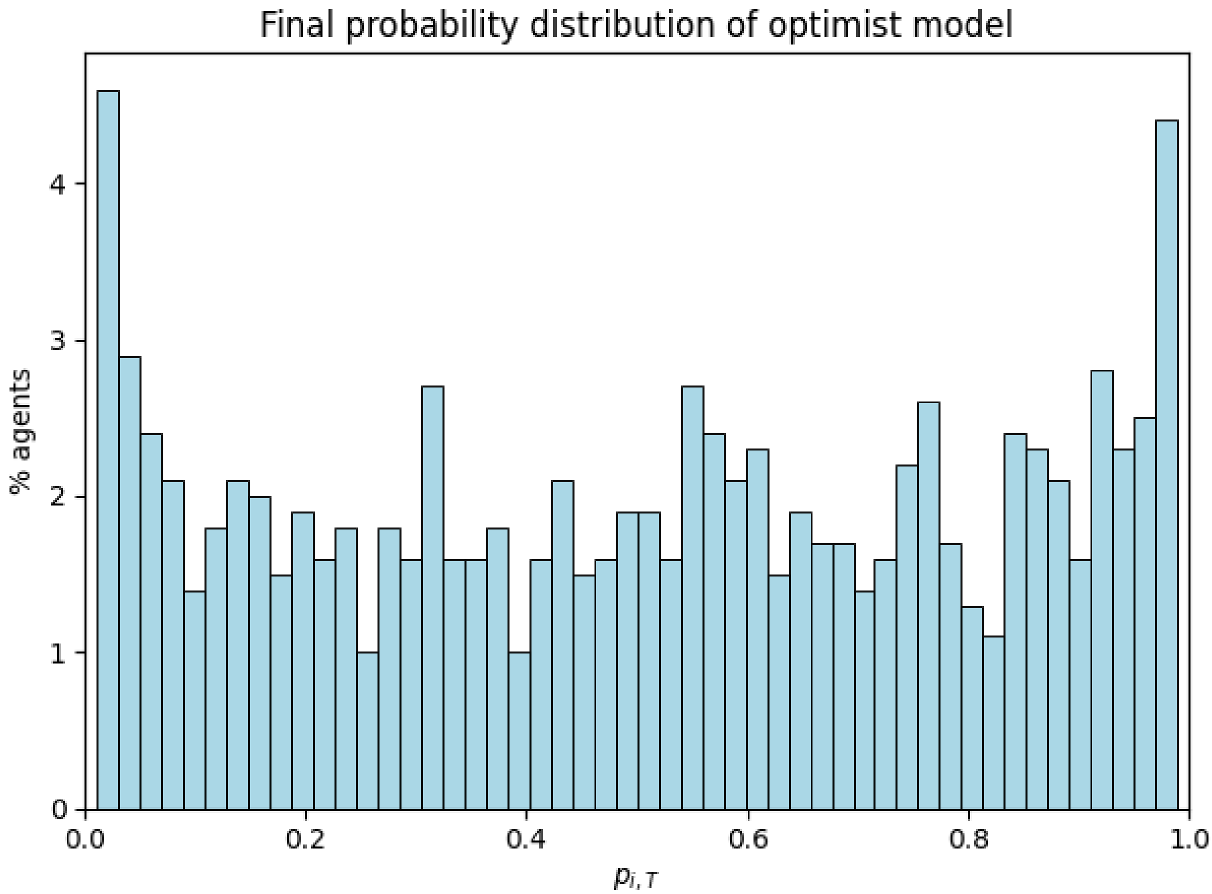

Figure 13 Subjective probability distribution at round t = 100 for a typical run of the model with heterogeneous agents (directional reasoning parameter α = 0.5). The probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model.

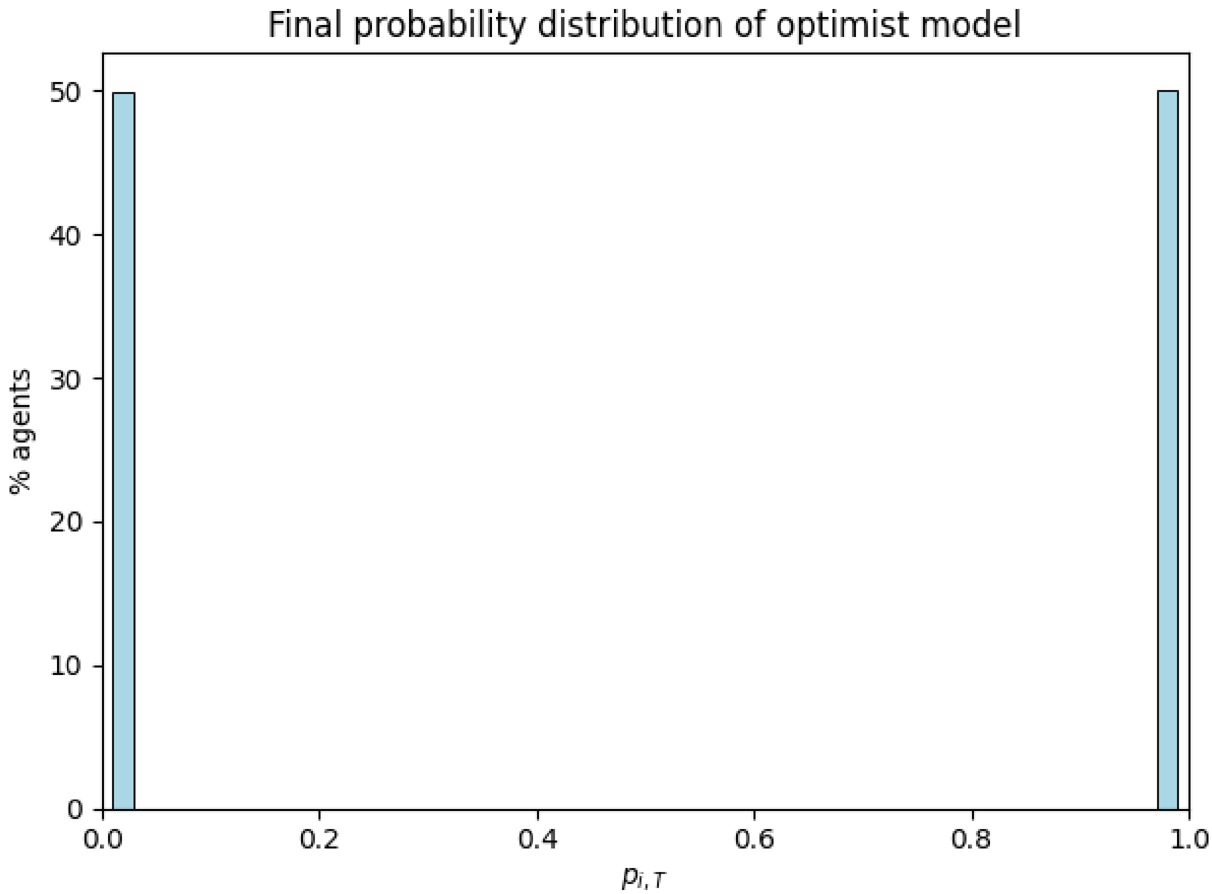

Figure 14 Subjective probability distribution at the steady state for a typical run of the model with heterogeneous agents (directional reasoning parameter $\alpha = 0.5$). The probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model.

One Lobbyist Scenario, homogeneous agents, $\lambda = 0.2$

In the following numerical exercise, the probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model. We have one lobbyist supporting the optimist model, and uniform agents with low underreaction.

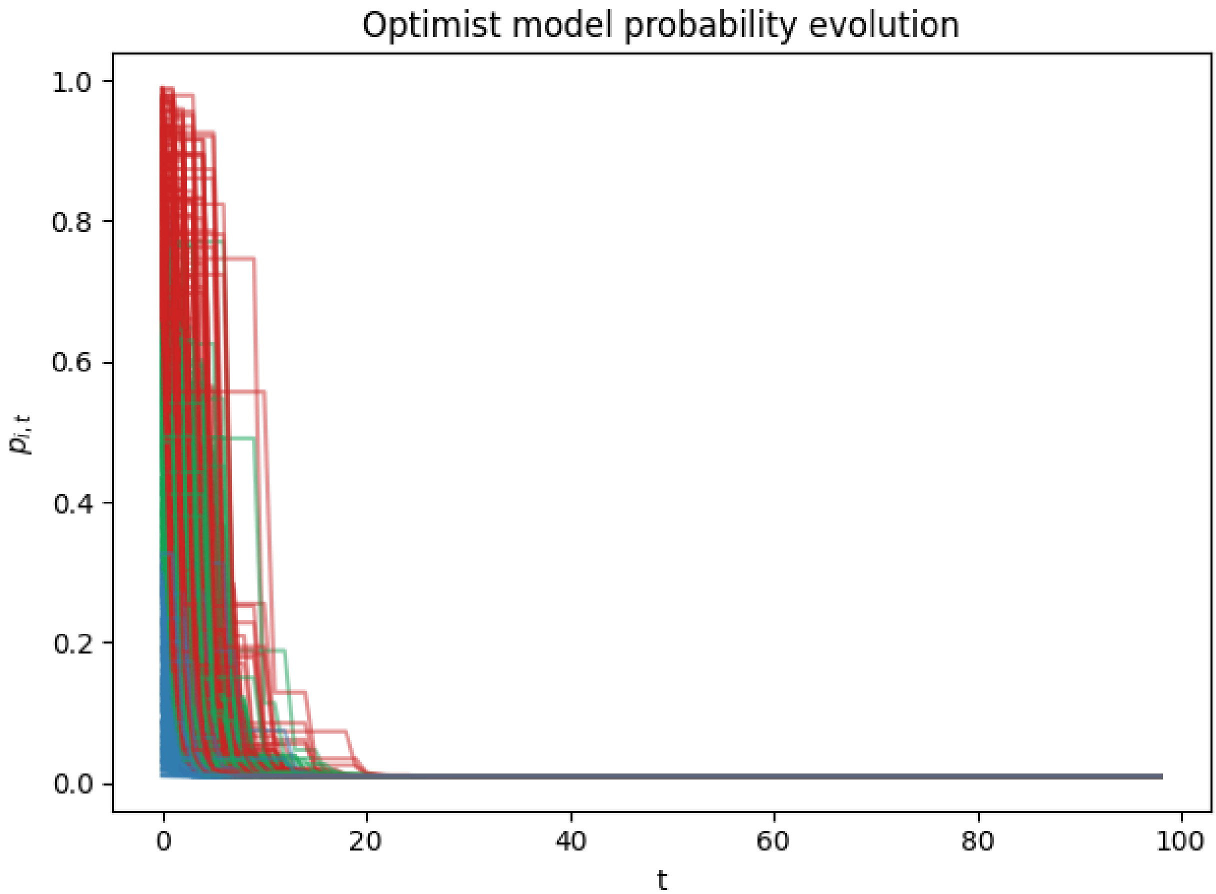

Figure 15 Evolution plot of agents' subjective probabilities for a typical run of the model with one lobbyist and homogeneous agents (underreaction parameter λ = 0.2). The probability of the negative climate event is p_p = 0.99 for the pessimist model and p_o = 0.01 for the optimist model.

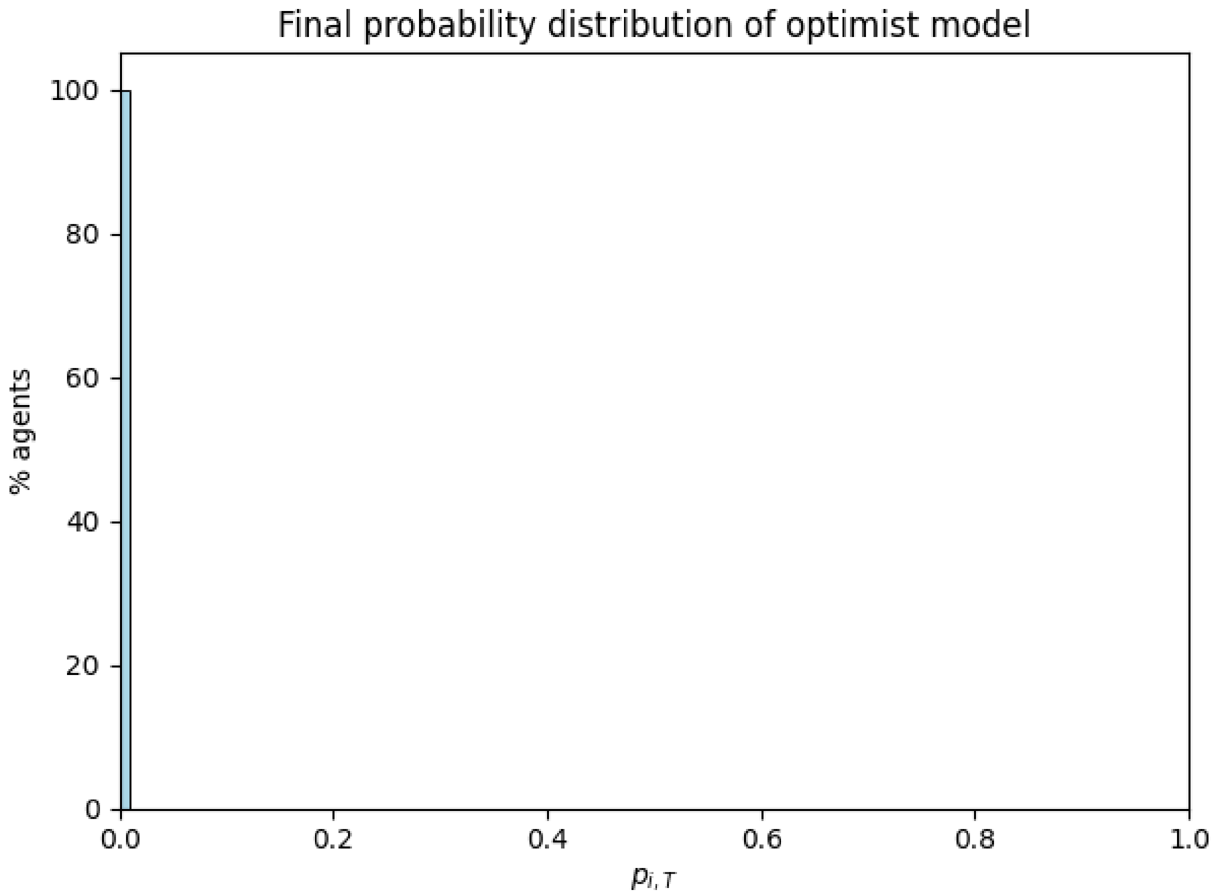

Figure 16 Subjective probability distribution at the steady state for a typical run of the model with one lobbyist and homogeneous agents (underreaction parameter $\lambda = 0.2$). The probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model.

One Lobbyist Scenario, heterogeneous agents, $\alpha = 0.5$

In the following numerical exercise, the probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model. We have one lobbyist supporting the optimist model and heterogeneous agents with directional reasoning.

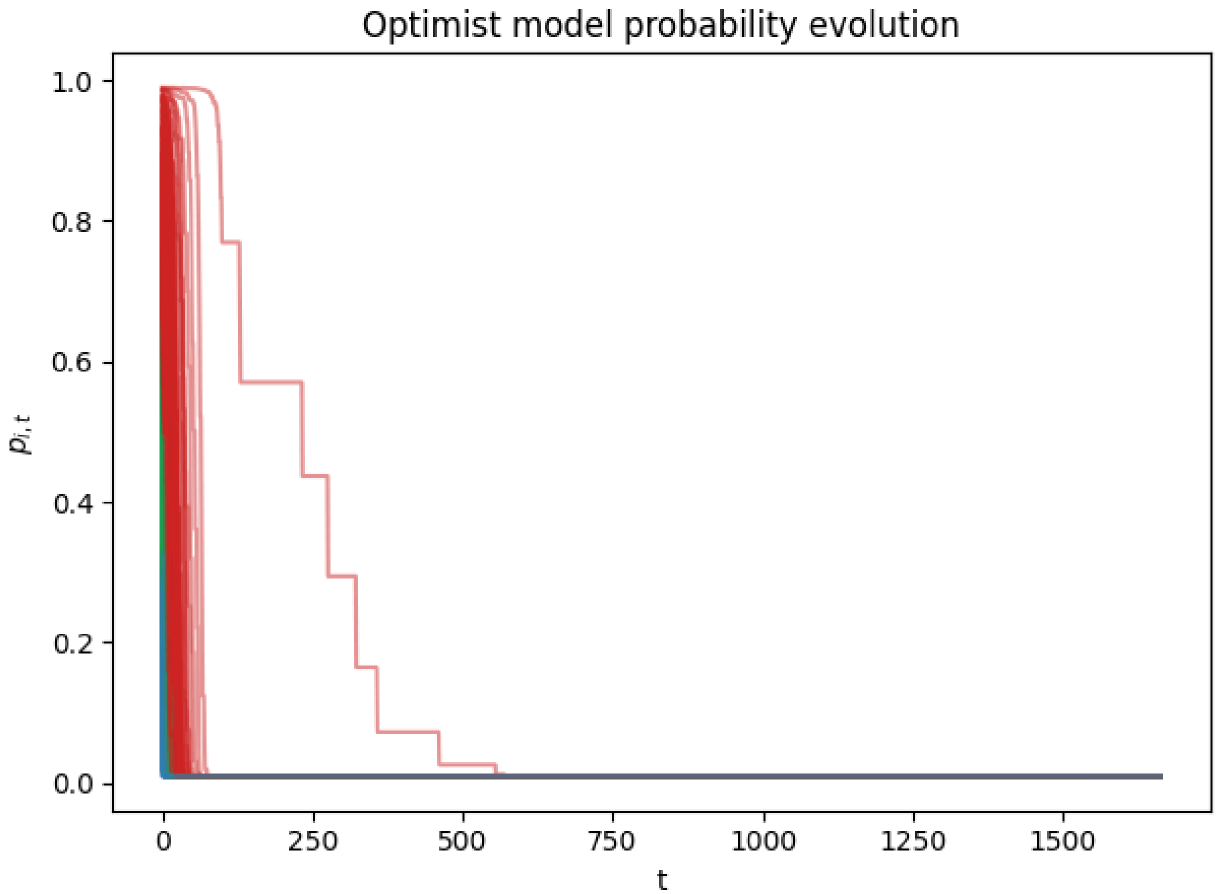

Figure 17 Evolution plot of agents' subjective probabilities for a typical run of the model with one lobbyist and heterogeneous agents (directional reasoning parameter α = 0.5). The probability of the negative climate event is p_p = 0.99 for the pessimist model and p_o = 0.01 for the optimist model.

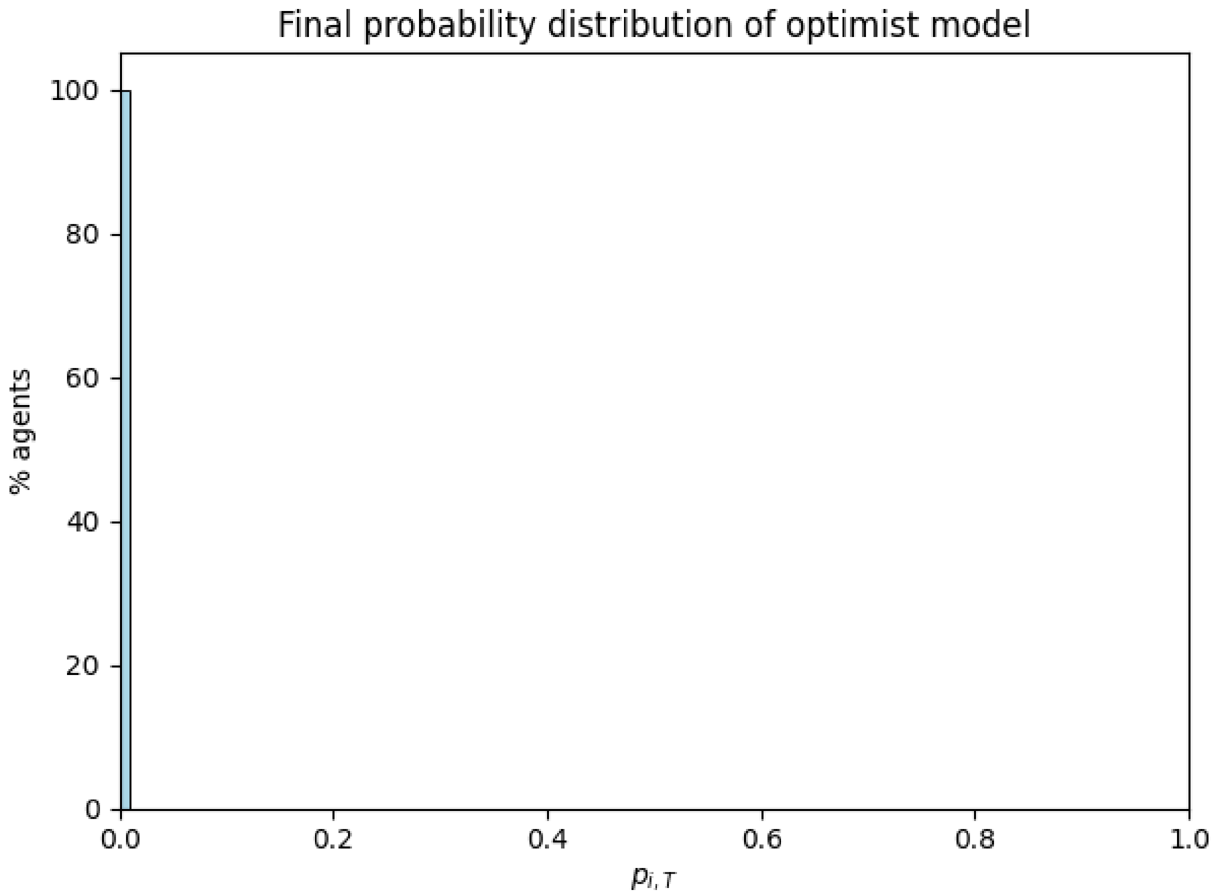

Figure 18 Subjective probability distribution at the steady state for a typical run of the model with one lobbyist and heterogeneous agents (directional reasoning parameter α = 0.5). The probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model.

Two lobbyists scenario, homogeneous agents, λ= 0.2

In the following numerical exercise, the probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model. We have two opposing lobbyists with the same strategy and unlimited budget.

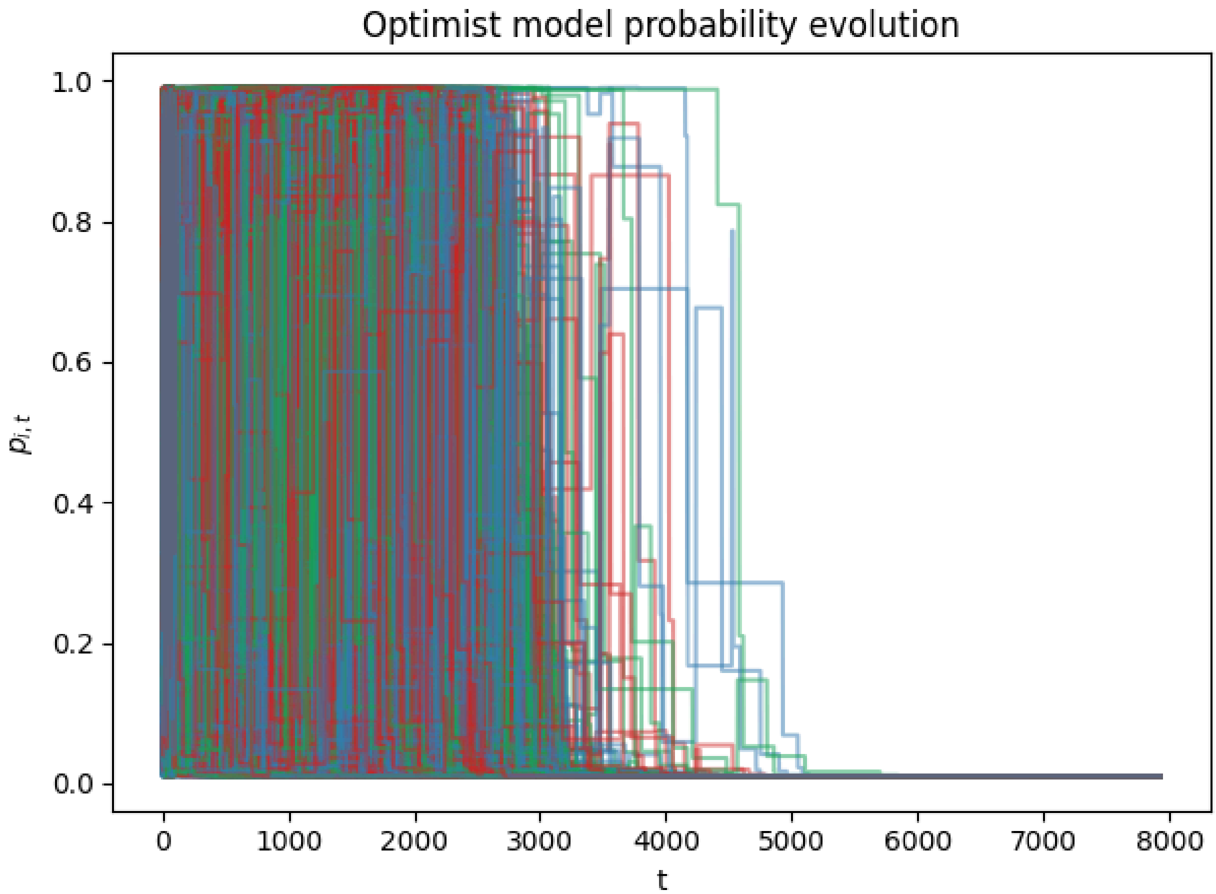

Figure 19 Evolution plot of agents' subjective probabilities for a typical run of the model with two lobbyists and homogeneous agents (underreaction parameter $\lambda = 0.2$). The probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model.

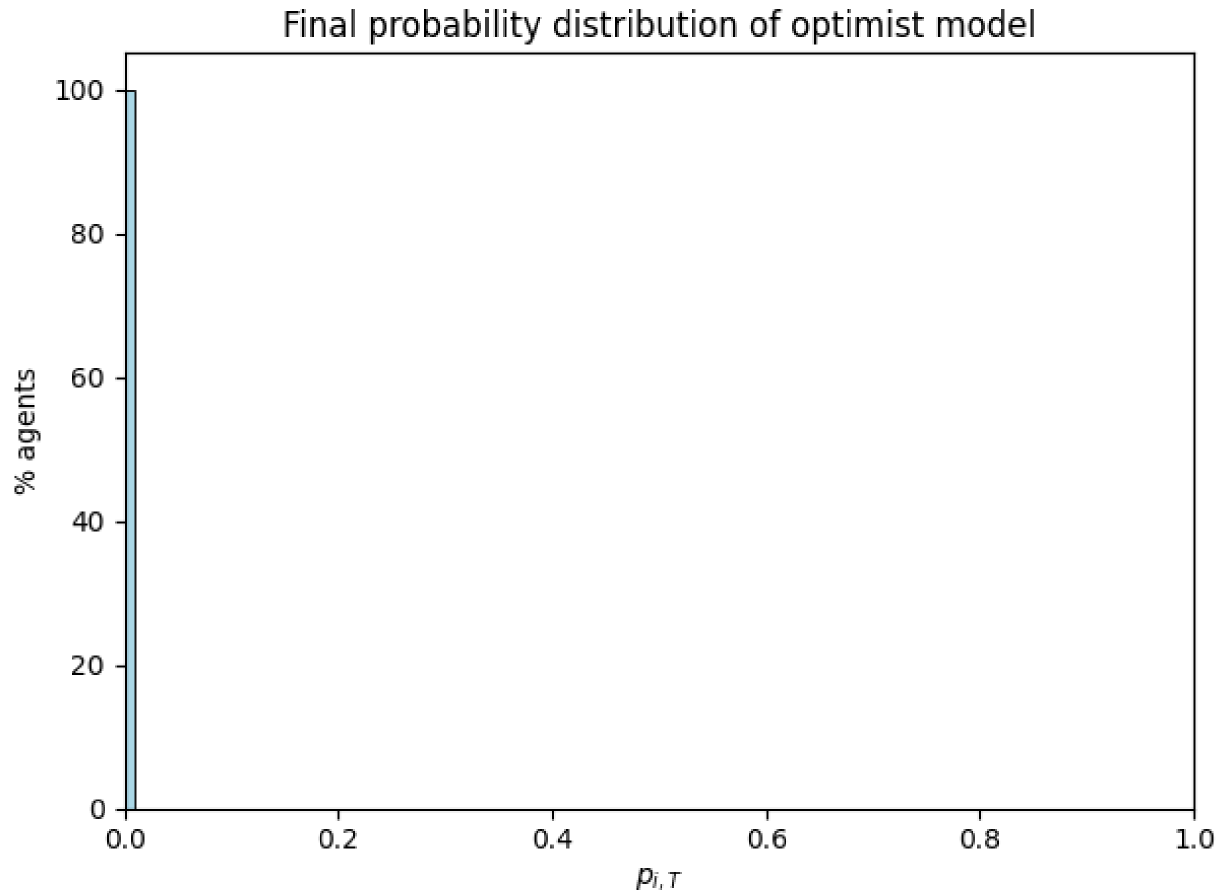

Figure 20 Subjective probability distribution at the steady state for a typical run of the model with two lobbyists and homogeneous agents (underreaction parameter λ = 0.2). The probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model.

Two lobbyists scenario, heterogeneous agents, α= 0.5

In the following numerical exercise, the probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model.

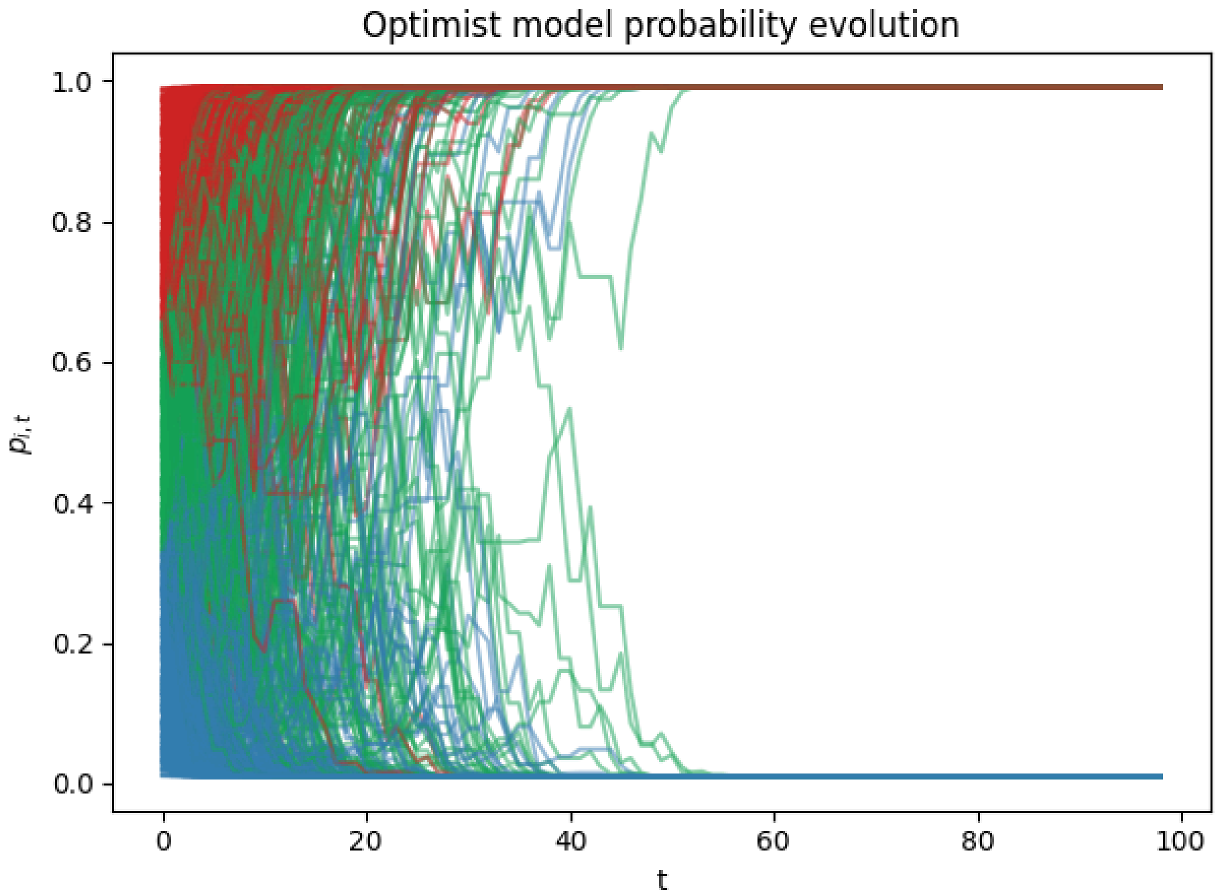

Figure 21 Evolution plot of agents' subjective probabilities for a typical run of the model with two lobbyists and heterogeneous agents, p_o = 0.01, p_p = 0.99, α = 0.5

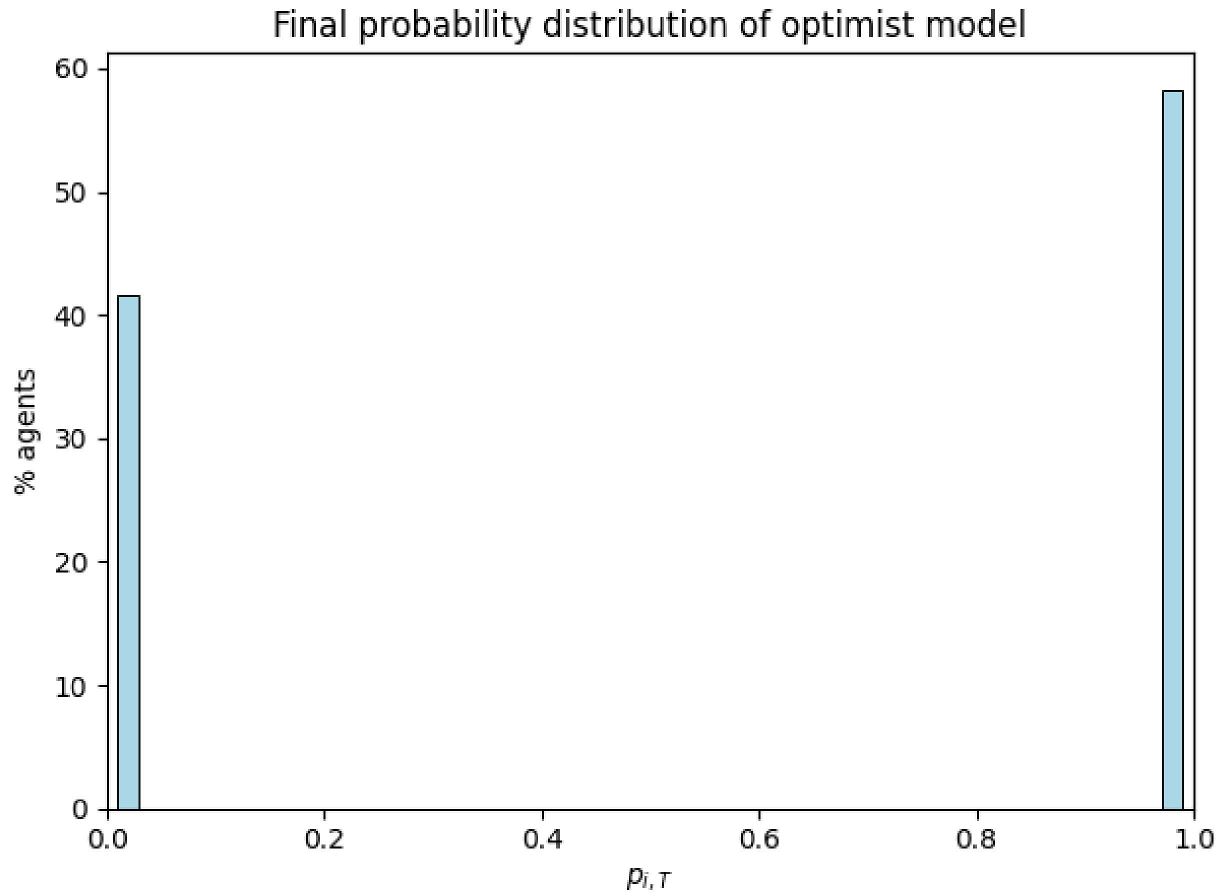

Figure 22 Subjective probability distribution at the steady state for a typical run of the model with two lobbyists and heterogeneous agents (directional reasoning parameter α = 0.5). The probability of the negative climate event is $p_p = 0.99$ for the pessimist model and $p_o = 0.01$ for the optimist model.

References

Bottazzi, G., Giachini, D., & Ottaviani, M. (2023). Market selection and learning under model misspecification. Journal of Economic Dynamics and Control, 156, 104739.

Botzen, W., Duijndam, S., & van Beukering, P. (2021). Lessons for climate policy from behavioral biases towards COVID-19 and climate change risks. World Development, 137, 105214.

Centola, D. and M. Macy (2007). Complex contagions and the weakness of long ties. American journal of Sociology 113 (3), 702–734.

Das, A., S. Gollapudi, and K. Munagala (2014). Modeling opinion dynamics in social networks. In Proceedings of the 7th ACM international conference on Web search and data mining, pp. 403–412.

Druckman, J. N., & McGrath, M. C. (2019). The evidence for motivated reasoning in climate change preference formation. Nature Climate Change, 9(2), 111-119.

Dunlap, R. E. and P. J. Jacques (2013). Climate change skepticism and denial: An introduction. American Behavioral Scientist 57 (6), 691–698.

Epstein, L. G., Noor, J., & Sandroni, A. (2010). Non-bayesian learning. The BE Journal of Theoretical Economics, 10(1), 0000102202193517041623.

Hegselmann, R., U. Krause, et al. (2002). Opinion dynamics and bounded confidence models, analysis, and simulation. Journal of artificial societies and social simulation 5 (3).

Li, T. and H. Zhu (2020). Effect of the media on the opinion dynamics in online social networks. Physica A: Statistical Mechanics and its Applications 551, 124117.7

Massari, F. (2020). Under-reaction: Irrational behavior or robust response to model misspecification?. Available at SSRN 3636136.

Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). Microeconomic theory. Oxford University Press.

Nisbet, M. and E. M. Markowitz (2014). Understanding public opinion in debates over biomedical research: looking beyond political partisanship to focus on beliefs about science and society. PloS one 9 (2), e88473.

Robinson, J., M. Bradley, P. Busby, D. Connor, A. Murray, B. Sampson, and W. Soper (2006). Climate change and sustainable development: Realizing the opportunity. AMBIO: A Journal of the Human Environment 35 (1), 2–8.

Rowell, A. (2013). Oil companies and climate change: Denial and delay. The Guardian 9.

Schlichting, I. (2013). Strategic framing of climate change by industry actors: A meta-analysis. Environmental Communication: A Journal of Nature and Culture 7 (4), 493–511.

Sirbu, A., V. Loreto, V. D. Servedio, and F. Tria (2013). Opinion dynamics with disagreement and modulated information. Journal of Statistical Physics 151, 218– 237.

Sirbu, A., V. Loreto, V. D. Servedio, and F. Tria (2017). Opinion dynamics: models, extensions and external effects. Participatory sensing, opinions and collective awareness, 363–401.

Sobkowicz, P. (2020). Whither now, opinion modelers? Frontiers in Physics 8, 587009.

Stokes, L. C. and C. Warshaw (2017). Renewable energy policy design and framing influence public support in the united states. Nature Energy 2 (8), 1–6. 8

Yeganeh, A. J., A. P. McCoy, and T. Schenk (2020). Determinants of climate change policy adoption: A meta-analysis. Urban Climate 31, 100547.