# Tolerance in Polarized Societies: A Game Theoretic Agent-Based Model

Eli Propp
The University of Waterloo
Waterloo, Canada
ehpropp@uwaterloo.ca

Nicholas Khorasani
The University of Waterloo
Waterloo, Canada
nckhoras@uwaterloo.ca

## ABSTRACT

This paper investigates the dynamics of ideological polarization within societies using game theoretic principles and agent-based simulations. We focus on modelling scenarios where individuals maintain steadfast beliefs while adapting their tolerance toward others. Through simulations, we explore the convergence of agent strategies and analyze the influence of varying parameters, such as simulation length and the initial strategies of the agents, on the emergence of cooperative behaviors. Our findings reveal that, despite societal polarization, a significant portion of agents tend to converge on high-tolerance strategies, especially when starting biases are not overwhelmingly skewed towards low tolerance. These results suggest that, even in polarized environments, a propensity exists towards cooperative engagement and inclusive dialogue when appropriately incentivized, highlighting the potential for fostering constructive interactions within diverse societies.

**ACM Reference Format:**
Eli Propp and Nicholas Khorasani. 2025. Tolerance in Polarized Societies: A Game Theoretic Agent-Based Model. In *Proceedings of ACM Conference (Conference'17).* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

## 1 INTRODUCTION

In recent years, democratic societies, including countries such as the United States and Canada, have experienced a concerning trend towards ideological polarization. This phenomenon is marked by a growing number of individuals who perceive societal divisions through a binary lens, categorizing others as either 'with' or 'against' their ideological beliefs, with little room for nuanced discourse or compromise. Consequently, the fabric of civil discourse becomes increasingly strained, and the prevalence of productive conversations on pressing societal issues diminishes.

As people's views on the world become more intertwined with their self-identification, they are less likely to change their opinions, as that would mean they, as a person, were previously mistaken. However, people with different views can still have productive conversations and tolerate people who oppose them. For example, many people of different religious backgrounds can hold diametrically opposing views yet still be friends.

To this end, this paper aims to develop a model to simulate what might lead people to become more tolerant of others over time. By simulating agent interactions where individuals maintain steadfast beliefs while having the flexibility to adjust their tolerance levels, our model mirrors real-world scenarios where ideological convictions remain static and attitudes towards cooperation and dialogue can evolve. This nuanced representation enables us to explore how changes in various parameters, discussed in section 3.2, impact the likelihood of constructive engagement and the emergence of cooperative strategies within polarized societies.

To accomplish this task, this paper leverages principles from the field of game theory to simulate agent interactions. Specifically, we leverage game theory to develop a simulation-based approach for investigating the dynamics of agent interactions within polarized societies. Agents in the simulation are self-interested, a core idea in game theoretic models, looking to maximize their utility as they interact with others. The agents will learn from their interactions to update their strategy. After enough simulation time has passed, we will investigate if there is any convergence on an optimal strategy to learn what might lead agents to prefer having a high tolerance towards other agents' beliefs. By modelling societal interactions as strategic games and exploring strategies for fostering productive dialogue, our research aims to contribute to the broader discourse on mitigating ideological polarization and promoting inclusive, participatory democracies.

The rest of this paper is structured as follows. Section 2 presents related works in this area. Section 3 describes the game theory model used in the simulation and the simulation environment. In section 4, we discuss our simulation findings. Section 5 has our concluding thoughts.

## 2 RELATED WORKS

Several studies have explored polarization dynamics in various contexts. However, most of them develop models and simulations where agents update their opinions instead of keeping them stagnant. Macy et al. [1] constructed an environment where agents possess modifiable opinion vectors that update based on interactions with others. Their study, focusing on polarization in an n-party system (simplified to two parties for simulations), examines how polarization spreads and at what threshold parties become too polarized to collaborate on resolving an exogenous event they initially agreed upon. They evaluated the impact of parameters such as the influence and timing of shocks on agent polarization. Notably, this work does not incorporate agent utility or decision-making, with


*Conference'17, July 2017, Washington, DC, USA*

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

agents updating opinions solely based on a predefined formula after each interaction.

Some works, like Rychwalska and Roszczyńska-Kurasińska in [2], investigate the causes leading to opinion polarization using agent-based models, highlighting the role of social media as a significant contributor. Others, such as Schweitzer, Grivachy, and Garcia in [3], develop agent-based models to examine how emotions influence the polarization of opinions.

A common theme among most works utilizing agent-based models to investigate polarization is their emphasis on causes or dynamics under the assumption that opinions will change. While this assumption holds true in many scenarios, our goal is to model a scenario where people's beliefs remain constant while their outlook towards others can change. We contend that such scenarios are increasingly prevalent in society, where polarization has evolved from a possibility to an established reality. Moreover, some existing works do not incorporate a game-theoretic approach. Therefore, we present a novel perspective on polarization dynamics by integrating game theory principles, offering unique insights into societal polarization dynamics.

# 3 METHODOLOGY

## 3.1 The Model

This project aims to simulate societal interactions in which agents maintain fixed beliefs while adapting their strategies to maximize utility. Each agent's strategy determines their interaction tolerance level, influencing their utility. The model and simulation framework are elaborated upon in the subsequent subsections. For this simulation, 1000 agents were utilized.

*3.1.1 Opinion Generation.* Agent opinions were generated using a bimodal distribution centred around 0. Opinions close to 0 represent a centrist perspective, and opinions with large absolute values represent the extreme perspectives.

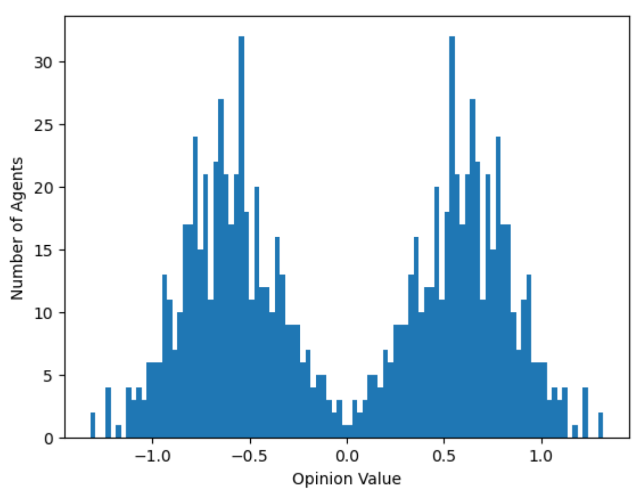


**Figure 1: Distribution of agent opinions**

Figure 1 shows the distribution of agent opinions used in the simulations. The distribution of agent opinions was kept consistent across all simulation runs, mitigating the potential bias from variations in opinion values and counts. The mean of each distribution is $\mu = \pm\, 0.6$, with a standard deviation $\sigma = 0.25$. The range of values for the opinions is $[-1.5, 1.5]$, the bounds of which each denotes either extreme. The total number of opinions generated was 1000; one opinion per agent.

In this model, agent opinions remain static. This choice reflects real-world scenarios involving polarizing topics, such as religion and politics, where opinions tend to resist change while tolerance levels may evolve.

*3.1.2 Strategies.* The set of actions available to each agent is $A = \{1, 2, 3\}$, where each $a \in A$ corresponds to a tolerance level. Tolerance represents an agent's willingness to engage in constructive dialogue. It can also be thought of as the amount of effort an agent is willing to put into having a productive conversation. A tolerance of 1 corresponds to low tolerance (i.e. a low willingness to put effort into a conversation), whereas a tolerance of 3 reflects a high willingness to invest effort.

Agents employ a mixed strategy according to the weights of each action, discussed in section 3.1.4.

*3.1.3 Utility Function.* The utility function for each agent is defined as:

$$u_i(a) = \text{softplus}\big(\max(a_i, a_{-i}) - d\big) \times \left(\frac{a_i + a_{-i}}{d} \cdot \text{reward} - \text{cost} \cdot a_i \cdot d\right) \quad (1)$$

Equation 1 was used for multiple reasons. First, it incorporates softplus acts to dampen utility when no discussion takes occurs, reflecting situations where agents fail to collaborate due to large opinion disparities. The softplus function is defined by

$$\text{softplus}(x) = \log(1 + e^x)$$

This is determined by looking at the difference between the agent with the higher tolerance chosen ($\max(a_i, a_{-i})$) and the distance ($d$) between the two agents opinions (interactions will be discussed more in section 3.1.4). If the distance is too large, the agents will be unable to collaborate, resulting in decreased utility.

The utility function's second component factors in reward and cost terms, modulated by agent effort and opinion distance to capture interaction dynamics. The variables *reward* and *cost* in equation 1 are both kept constant at 1. The reward is proportional to the sum of the efforts of each agent, $a_i + a_{-i}$, and inversely proportional to the distance, $d$. However, the cost is proportional to both agent $i$'s effort $a_i$ and the distance. When two people meet, their combined effort will result in a higher reward, while being further from the other will result in a lower reward. But when looking at the cost, the more effort an agent puts in, the higher their cost, on top of a higher cost for having a higher distance.

Notably, the utility function avoids strict dominance of high tolerance strategies. Otherwise, there would be an obvious choice that should rapidly converge for the agent: always choose High tolerance (i.e. $a_i = 3$ with probability 1). This property makes the interactions and chosen strategy more interesting for analysis, while keeping the the utility function similar to a real world scenario: putting in too much work can cost you more in certain interactions.

For example, take two agents, $A_1$ and $A_2$, whose opinions have a distance of 1. When $A_1$ and $A_2$ both choose $a = 1$, $u_1 = 0.69$, and if $A_2$ chose a tolerance of $a_1 = 2$, while $A_2$ keeps $a_2 = 1$, $u_1$ would increase to 1.31. However, is the distance was 2, when $a = 1$, $u_1 = -0.31$, and when $a_1 = 2$ and $a_2 = 1$, $u_1$ would decrease to -1.73.

*3.1.4 Learning.* No-regret learning was employed to facilitate agents updating their strategy. No-regret learning models look at the outcome of each possible action for each interaction and update the probabilities of choosing each action based on how they would have performed in that interaction. The specific algorithm chosen to implement no-regret learning was the multiplicative weights (MW) algorithm.

The MW algorithm was chosen for the following properties:

- MW can handle $N$ distinct actions.
- MW can handle arbitrary and continuous costs/loss in $[0, 1]$.

The general structure of the MW algorithm, described in [4], is as follows,:

(1) In round $1, ..., T$, the algorithm choose some expert $i^t$
(2) Each expert $i$ experiences loss $l_i^t \in [0, 1]$
(3) Algorithm experiences the loss of the expert it chooses $l_A^t = l_{i^t}^t$
(4) Total loss of expert $i$ is $L_i^T = \sum_{t=1}^{T} l_i^t$
(5) Total loss of the algorithm is $L_A^T = \sum_{t=1}^{T} l_A^t$

The above structure allows the loss of each agent to be stored for later analysis. The pseudo-code of the basic MW algorithm, described in [4], is:

**Algorithm 1** MW Algorithm

```
Set weights w_i^t to 1 for all actions i
for t = 1 to T do
    Let W^t = Σ_{i=1}^N w_i^t
    Choose action i with probability w_i^t / W^T
    For each i, set w_i^{t+1} ← w_i^t · exp(−ε l_i^t)
end for
```

Note: The first step, where the weights are set, was not the same for all simulations. This will be discussed further in section 3.2.3.

The loss is defined as

$$l_i^t = 1 - \frac{u(a_i) - min(u(a))}{max(u(a)) - min(u(a))}$$

The loss must be normalized to a range of $[0, 1]$ as described above. Instead of directly using the normalized value, we subtract it from 1. This was done when the fractional part was 1, which means the agent's chosen strategy yielded the maximum utility for that interaction, which would result in a loss of 0. In this implementation, $\epsilon = 0.1$.

## 3.2 Simulation

Given the model above, the simulation serves to analyze the impact of varying parameters on agent strategy profiles, which are simply the probability of each agent choosing each action.

The remainder of this section will outline the general simulation and interaction structure and discuss the different parameters that were changed for each simulation.

*3.2.1 Simulation Structure.* Each simulation involves a total of $N = 1000$ agents and runs for $T$ time steps. In each time step, agents sequentially interact with randomly chosen counterparts. That is, agent 1 chooses a random agent to interact with, determines its strategy based on its current weights, and the current weights are stored for further analysis. Then, the utility is determined, and weights are updated as described in the MW algorithm in section 3.1.4. Then Agents 2 through N then go through the same process. The time step is concluded once each agent has experienced exactly one interaction.

*3.2.2 Number of Time Steps.* The first parameter modified between simulation runs was the length of the simulation. The goal of varying the simulation length is to observe trends in the convergence of strategy profiles. Since the distance between agent $i$'s opinion and the agent it chooses is part of the utility function in equation 1, there may be trends in converged strategy profiles based on the agent's opinion if the strategy profile converges.

*3.2.3 Starting Strategy.* The other parameter that changed between simulations was the starting strategy profile of the agent. At first, all agents started with all weights for each action set to 1, as described in algorithm 1. This means each agent starts with equal probabilities of choosing each tolerance. However, in society, not everyone has the same level of tolerance. More diverse initial strategy profiles are explored to emulate real-world variance in tolerance levels. Simulation runs commence with agents exhibiting varied preferences for low, medium, or high tolerance. An example of this is could be 80% starting with $[4, 1, 1]$, 15% with $[1, 4, 1]$, and 5% starting with $[1, 1, 4]$, a low, medium, or high preference of 66.7%, respectively. As agents can start with different probabilities for their initial strategies, this may influence what the optimal strategy for any agent will be.

This parameter variation aims to identify thresholds wherein optimal high-tolerance strategies prevail, considering fixed initial opinions and potential opinion-strategy correlations. As mentioned before, since distance is incorporated in the utility function in equation 1, and agents starting opinion is fixed, agents with different opinions may end with different optimal strategies.

# 4 EXPERIMENTAL RESULTS

## 4.1 Strategy Convergence With Equal Probability Starting Strategies

This section investigates the convergence of agent strategy profiles when starting with equal probability strategies. The point of interest is to see if a high tolerance strategy is preferred, given it is not a strictly dominant strategy. The initial simulations ran with each agent having no preference to any action, i.e. an equal probability for each. The convergence of pure strategies can be seen in figure 2.

Figure 2 shows the number of agents that hold pure and mixed strategies at each time step. The green line shows the number of agents playing a pure strategy of high tolerance, blue medium, orange low, and grey mixed. Figure 2 represents agents with opinions greater than 0, which is why the total number of agents is 500. The figure for agents with opinions less than 0 is the same.

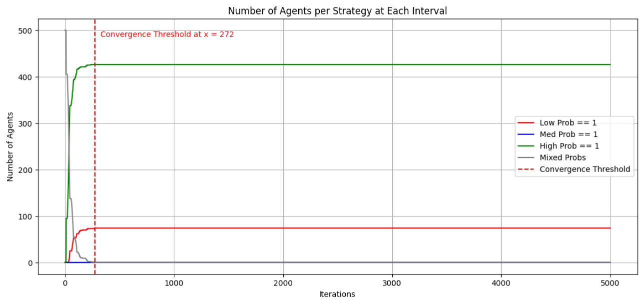


**Figure 2: Number of agents preferring each strategy over time**

At the start all agents play a mixed strategy with equal weights, shown by the grey line at 500 for $t = 0$. However, as interactions unfold, a significant portion converges to pure strategies, particularly towards high tolerance, as depicted by the green line in the figure. The number of agents playing a mixed strategy quickly drops 0, while the number of agents playing pure strategies of low or high tolerance converges at around the same time, approximately $t = 272$.

While all agents do not converge on a pure strategy of high tolerance, approximately 84% do. Figure 3 provides some insight as to why this is the case.

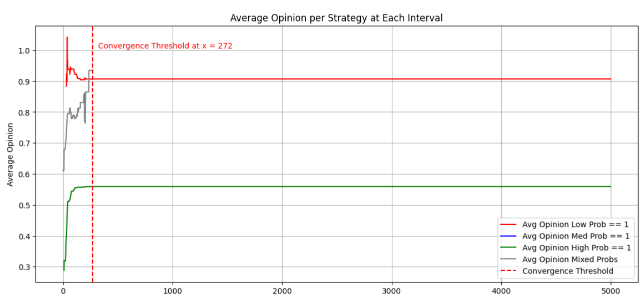


**Figure 3: Average opinion of agents playing each strategy over time**

Figure 3 shows the average opinion values for the same strategies from figure 2. As with the number of agents in figure 2, the lines for each strategy converge. The average opinion for agents with a pure strategy of low tolerance is around 0.91, and of high tolerance is 0.56 (it is the negative of each when plotting the agents with opinions below 0). Agents with more extreme opinions tend to favour low-tolerance strategies, while those with moderate opinions lean towards high tolerance. This aligns with the utility function in equation 1, where agents with extreme views experience lower utility due to greater average distance from others, incentivizing low tolerance. At a certain level, putting in more effort with a high tolerance strategy does not yield a higher utility based on their opinion.

## 4.2 Strategy Convergence With Biased Starting Strategies

Next, we studied the system when agents were initialized with imbalanced strategies. Figure 4 shows that even when 60% of the agents started with an 80% bias towards playing low tolerance, the majority of the agents in equilibrium played high tolerance. An 80% bias towards low tolerance means the starting weight vector for those agents is $[8, 1, 1]$, with the weights corresponding to low, medium, and high tolerances, respectively.

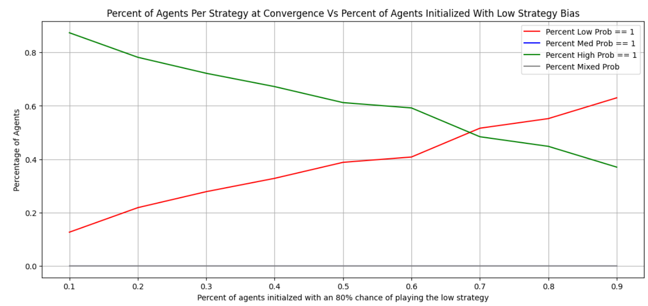


**Figure 4: Number of agents playing each strategy in equilibrium vs the biased starting state**

The crossover point in figure 4, where the same proportion of agents prefer a pure strategy of low and high tolerance, is around 68%. This means that when 68% of the agents start with an 80% bias towards low tolerance, the simulation converges to an even split of agents preferring low and high tolerance pure strategies.

We can also see that the average opinions of the agents playing each strategy in equilibrium changed as we changed the starting strategies in figure 5.

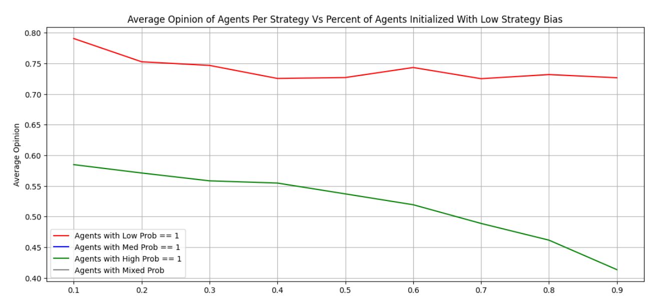


**Figure 5: Average opinion of agents playing each strategy in equilibrium vs the biased starting state**

As illustrated in Figure 5, altering the starting strategies influences the average opinions of agents at equilibrium. Taking step back, this is saying that for those agents playing a pure low strategy, the average opinion does not need to be as extreme as more agents start with a similar bias as you. For those playing a pure high strategy, it shows that the average opinion will become more neutral, as the agents whose opinions were at the bounds of where it made sense to play a pure high strategy will slowly become attracted to playing a pure low strategy.

## 4.3 Agent Loss

Plotting the agent loss based on their learned strategies provides some interesting insights. The goal was to examine the loss of the strategies employed by agents, to see which benefited the most.

*4.3.1 Equal Probability Starting Strategies.* Figures 6 and 7 showcase the average loss relative to agent opinions after 500 and 10000 iterations, respectively. The blue points are agents that converged on a pure high strategy, the red on a pure low strategy.

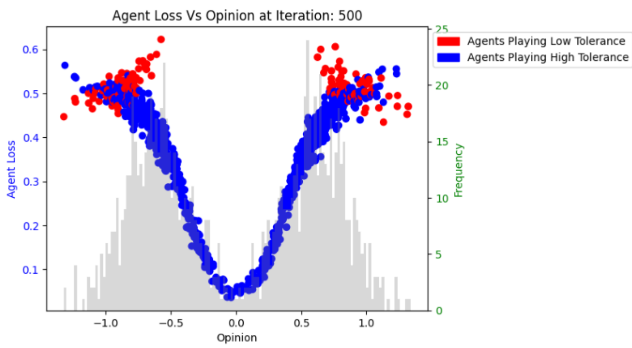


**Figure 6: Average agent loss and frequency relative to opinion after 500 iterations**

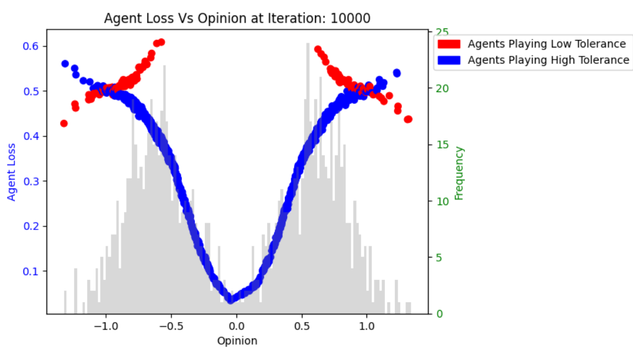


**Figure 7: Average agent loss and frequency relative to opinion after 10000 iterations**

Both figures 6 and 7 show the same trend, with 7 being a smoother representation as it was a longer simulation.

High tolerance agents experience increased loss with extreme opinions, whereas low tolerance agents benefit from extreme views. Interestingly, the loss of the low tolerance agents intersects with the high tolerance agents around and opinion of ±1, where playing low tolerance with an even more extreme opinion would result in a lower average loss than playing high tolerance, whereas with a less extreme opinion it would result in a higher average loss.

*4.3.2 Biased Starting Strategies.* When the starting strategies of the agents are biased, there is a similar trend in the loss as above, but with additional loss some cases. Figures 8 and 9 show the average loss plots, but with 10% and 80% biased towards a low tolerance strategy at the start, respectively.

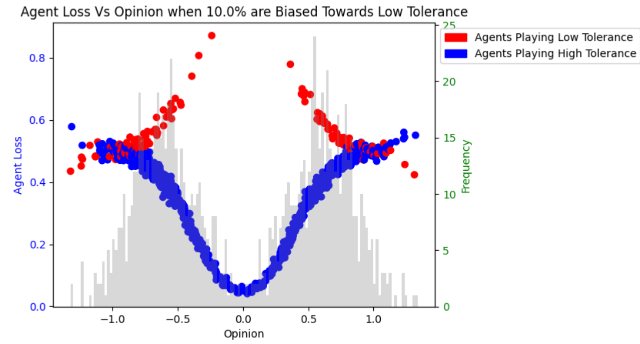


**Figure 8: Average agent loss and frequency relative to opinion with 10% biased to low tolerance**

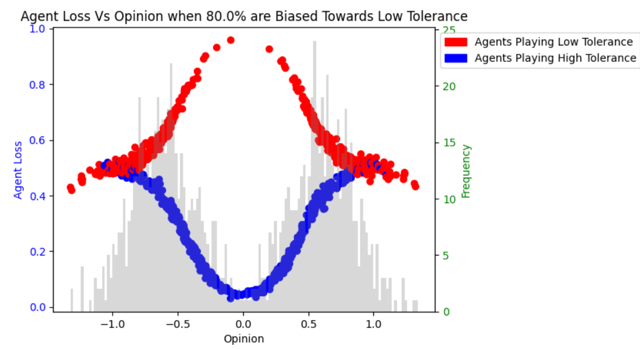


**Figure 9: Average agent loss and frequency relative to opinion with 80% biased to low tolerance**

When the 10% of the agents are biased towards a low tolerance strategy as in figure 8, we see the same trend as in section 4.3.1: high tolerance agents have increased loss at their opinions are more extreme, low tolerance agents have a decreased loss. However, there is a slight difference between the loss for the low tolerance agents compared to an equal-weighted starting strategy as in section 4.3.1. The peak average loss for low tolerance agents in this scenario trends towards 1, as more agents with opinions closer to 0 end up playing a pure low tolerance strategy. This is even more pronounced in figure 9 where 80% of the agents start with a bias towards a low tolerance strategy. This is not the same for agents that play a high tolerance strategy in all scenarios, whose average loss reaches a max of approximately 0.55 with opinions more extreme than ±1.

Figure 10 shows the average loss of agents playing low or high pure strategies based on the proportion of agents that start with a bias to low tolerance.

When looking at the average loss in figure 10, signified by the solid blue a red lines with dots (low and high tolerance, respectively), we can see a slight increase for agents playing a low tolerance strategy as the percentage of agents who start with a bias towards low tolerance increases. Conversely, there is a slight decrease in the average loss across the axis for those playing a high tolerance strategy. To understand why, we can look at the number of agents that are playing each of these strategies, shown with the dashed

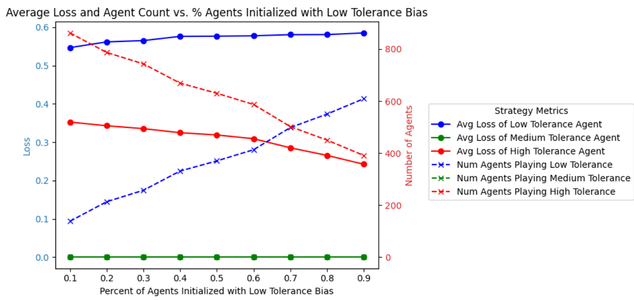


**Figure 10: Average loss of low and high pure strategies vs the percentage of agents initialized with a bias for low tolerance**

lines with x's with the same colour breakdown. Similar to what we saw in section 4.2, specifically figures 4 and 5, as more agents start biased to a low tolerance strategy, more converge on a pure low tolerance strategy, which brought the average opinion of these agents down. The same is true for the agents that converge on a high tolerance strategy, but less agents end up playing this strategy, and the average opinions of the ones that do are increasingly closer to 0.

## 5 CONCLUSION

The results of these simulations provide some interesting insights into the optimal tolerance strategies relative to agents' opinions. The first is that agents rarely, if ever, converged on playing a pure medium tolerance strategy. It was always better to put in as much effort as possible by playing a high tolerance or putting in a minimal amount of effort by playing a low tolerance. On top of this, agents rarely converged on a mixed strategy.

While we kept agent opinions static in these simulations to simulate topics where people rarely change their opinions, it is apparent that agents with the least extreme opinions and who converged on a pure high-tolerance strategy had the lowest loss and, therefore, performed the best.

When the starting weights were biased towards a low tolerance strategy, we found that when less than 70% of the agents had this low tolerance bias, a majority of the agents still converged on playing a pure high tolerance strategy, as shown in Figure 10. That is to say, even if over half of the agents start with a strong low tolerance bias, a majority of the agents still prefer a high tolerance. This is despite high tolerance not being a strictly dominant strategy, as discussed in section 3.1.3. However, agents that start with a bias towards low tolerance are more likely to end up converging on a low tolerance strategy, highlighting the importance of starting biases.

The results of these simulations give a promising outlook on society. While no simulation will perfectly match the intricate nuance of societal interactions, our simulation, which keeps opinions static but allows agents to become more (or less) tolerant of others, seems to capture the trend of polarization we see. In many scenarios, except when a large majority of agents start with an extremely low tolerance, most agents still prefer to be highly tolerant, even with more extreme opinions.

## REFERENCES

[1] Michael W Macy, Manqing Ma, Daniel R Tabin, Jianxi Gao, and Boleslaw K Szymanski. 2021. Polarization and tipping points. *Proceedings of the National Academy of Sciences* 118, 50 (2021), e2102144118.
[2] Agnieszka Rychwalska and Magdalena Roszczyńska-Kurasińska. 2018. Polarization on social media: when group dynamics leads to societal divides. (2018).
[3] Frank Schweitzer, Tamas Krivachy, and David Garcia. 2020. An agent-based model of opinion polarization driven by emotions. *Complexity* 2020 (2020), 1–11.
[4] Seyed M Zahedi. 2024. Learning in Games. Lecture.