

The social dilemma of online segregation: A dynamical model of platform choice

Rationality and Society

2026, Vol. 0(0) 1–47

© The Author(s) 2026



Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/10434631261428850

journals.sagepub.com/home/rss

Sven Banisch¹ , Dennis Jacob² , Tom Willaert³  and Ekehard Olbrich⁴ 

Abstract

To understand how social media shapes online discourse or contributes to polarization, we need models of collective online choice that link users' behavioral adaptation to the emergence of complex and dynamic digital environments. This study develops a dynamic model of platform selection based on Social Feedback Theory, using multi-agent reinforcement learning to capture how user decisions are shaped by past rewards across platforms. A key parameter (μ) governs whether users seek approval from like-minded peers or exposure to opposing views. Agent-based simulations combined with dynamical-systems analysis reveal a social dilemma: even when users value diversity, collective dynamics can trap online environments in polarized echo chambers that reduce overall user satisfaction. Above a critical diversity threshold, a different equilibrium appears in which one large, integrated platform dominates while smaller platforms persist as extremist niches. In an

¹Institute of Technology Futures, Department of Sociology and Computational Sociology, Karlsruhe Institute of Technology, Karlsruhe, Germany

²EECS Computer Science Division, UC Berkeley, Berkeley, CA, USA

³Brussels School of Governance, Vrije Universiteit Brussel, Brussels, Belgium

⁴Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany

Corresponding author:

Sven Banisch, Institute of Technology Futures, Department of Sociology and Computational Sociology, Karlsruhe Institute of Technology, Douglasstr. 24, Karlsruhe 76133, Germany.

Email: sven.banisch@universecity.de

intermediate regime, the two outcomes coexist, generating path dependence and hysteresis. We further demonstrate how modest, strategically targeted interventions—such as rewarding minority participation—can destabilize polarization and promote integrated discourse. The model shifts attention from belief change to participatory choice. It links micro-level learning dynamics to macro-level online fragmentation and informs mechanism-based interventions in the digital public sphere.

Keywords

social media, mathematical modeling, social dilemmas, online behavior, online segregation, multi-agent learning, computational social science

Introduction

Digital platforms have a profound impact on society, influencing how people communicate, access information, and express opinions (Boyd, 2014). While they have become integral to everyday life, they also raise concerns about polarization, misinformation, and the fragmentation of the public sphere (Sunstein, 2018; Tucker et al., 2018). Recent events surrounding platforms such as Twitter (now X) and the expansion of conglomerates like Meta—encompassing Facebook, Instagram, WhatsApp, and Threads—have intensified public scrutiny of the algorithms and governance structures that shape online discourse. Users increasingly respond to this environment by selecting platforms that align with their communicative preferences and values (Brady et al., 2023).

One of the primary concerns surrounding digital platforms is their role in reinforcing polarization and creating echo chambers, where users predominantly engage with like-minded individuals (Sunstein, 2018). This selective exposure not only amplifies existing opinions but also fosters ideological divides and undermines cross-cutting discourse (Stroud, 2010). A key driver of this phenomenon lies in the cultural norms and values embedded within each platform, which shape the kinds of opinions that are promoted or marginalized. Platforms differ in their normative structures—ranging from rigid moderation policies to more open, decentralized environments—resulting in distinct platform cultures (Scharlach and Hallinan, 2023; Van Raemdonck and Pierson, 2022). Users actively navigate these environments, often selecting online spaces that align with their own communicative preferences, values, and even political identities (Garrett, 2009; Knobloch-Westerwick, 2015). These preferences may reflect civic goals—such as

public debate and information access—or more interpersonal ones, including entertainment, affiliation, or self-expression (Lorenz-Spreen et al., 2024; Willaert and Olbrich, 2024). And while a user’s behavior may adapt, such preferences and identities tend to remain relatively stable over time. Simultaneously, the fluidity of the platform ecosystem—with frequent entry of new services and shifting platform policies—offers users continuous opportunities to switch when technical affordances or cultural dynamics no longer suit their needs. This raises important questions about how user choices, shaped by social feedback and platform-specific norms, contribute to the formation of either polarized spaces or diverse, open environments.

In particular, by means of their specific technical and social affordances, social media platforms might support key functions of the public sphere in liberal democracies. They might (1) facilitate the circulation of information and provide public access to information, (2) facilitate rational-critical debate among the public, (3) foster the creation and expression of collective identities, and (4) allow actors to coordinate and collectively perform actions (Willaert and Olbrich, 2024). When choosing which platforms to actively engage with, users might thus seek out those social media that best support the kind of interactions they value most highly. A user interested in belonging to a specific group or collective, might for instance look for those platforms that host like-minded individuals. On the other hand, users who wish to challenge their assumptions and expose themselves to new information, or wish to engage in rational debate with others, might be drawn to platforms that offer a diversity of perspectives and opinions. This dynamic selection process highlights the interplay between user motivations and platform affordances, underscoring how individual choices can shape the broader online discourse and, potentially, the structure of the digital public sphere.

To explain how either fragmented or integrated online spaces emerge from individual choices, this paper develops a theoretical model in which users can choose to express their opinion in a set of online communities called platforms. Drawing on agent-based simulations and rigorous bifurcation analysis, the paper reveals a social dilemma: *even if all users prefer to engage with a diversity of opinions, individual adaptation to social feedback may collectively lock the system into a suboptimal state of platform segregation composed of strongly opinionated echo-chambers.*

Our agent-based model (ABM) is rooted in social feedback theory (SFT) which draws on Homan’s ideas regarding social approval as a central mechanism of behavioral reinforcement (Homans, 1974). SFT transfers these principles to the context of online opinion dynamics (Banisch and Olbrich, 2019). The core idea is that agents play repeated communication games (Banisch et al., 2022; Gaisbauer et al., 2020) where they express their

opinions using a discrete set of options available on a social media platform (e.g. tweet/retweet on X or submission/comment on Reddit). Their choice propensities depend on the social feedback (e.g. likes/upvotes/downvotes) received in previous rounds of the collective game specified with the ABM.

Drawing on backward-looking reinforcement learning (RL), SFT adopts a bounded, procedural notion of rationality (Simon, 1978), where agents adapt their optimal strategy based on past rewards (Sutton and Barto, 2018). Hence, as opposed to discrete-choice models of critical transitions and tipping points (Granovetter, 1978; Kuran, 1989; Lopez-Pintado and Watts, 2008), RL-based models capture the adaptive nature of choice preferences, allowing us to study when and how different equilibria are reached out of non-equilibrium states through tacit coordination (Macy, 1991; Macy and Flache, 2002). In contrast to forward-looking models that rely on strong assumptions about cognitive capacity and foresight, independent RL (Matignon et al., 2012) makes modest assumptions about agents' information processing capacities, aligning with Thorndike's law of effect (Thorndike, 1911) as well as with more recent empirical research in social neuroscience (Izuma et al., 2010; Ruff and Fehr, 2014). From a sociological perspective, everyday interaction can be understood as a coordination game sustained by reinforcement dynamics (Vollmer, 2013), suggesting that RL captures core features of social life beyond abstract games. At the same time, RL provides a formal link between ABMs and equilibrium analysis (Fudenberg and Levine, 1998; Gaisbauer et al., 2020; Shoham and Leyton-Brown, 2008). This makes RL particularly relevant for modeling online interaction, where users repeatedly adapt their behavior to platform-specific feedback signals (Lindström et al., 2021), and where bounded, backward-looking adaptation may be a more realistic account of behavior than forward-looking optimization.

In this paper, we extend previous work on modeling online participation (Banisch et al., 2022; Gaisbauer et al., 2020; Törnberg et al., 2021). The model simulates a population of users, each holding one of two opposing opinions, who choose between a set of platforms for expressing their views. Users are influenced by two primary rewards: social approval, which is derived from interacting with like-minded individuals, and diversity, which comes from engaging with opposing views. Each platform is characterized by the proportion of users holding each opinion and the level of activity on the platform, which in turn affects user satisfaction. A key parameter μ governs the relative importance users place on social approval versus diversity. Through RL, users adapt their platform choices over time, updating their platform preferences based on the rewards they experience. This leads to dynamic changes in platform composition as users gravitate toward

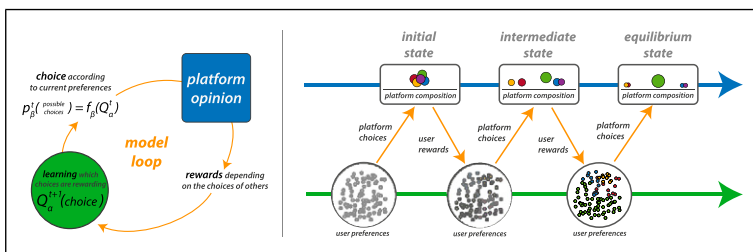


Figure 1. Platform choice models use reinforcement learning to address the co-evolutionary process of users' preferences and decisions to engage with different platforms, and the resulting platform ecology in which these decisions take place.

platforms that maximize their satisfaction. This approach hence captures the dynamic, evolving nature of online environments, where user preferences for different platforms co-evolve with the composition of the platform ecology (see Figure 1 for an illustration).

The analysis reveals two surprising equilibria emerging from the model as we increase diversity-seeking motivations (μ). First, the adaptive choices can lock the system into stable but suboptimal equilibria of platform segregation. Even when users value diversity higher than homophily, all platforms settle into polarized echo chambers, where users cluster around like-minded peers. Second, under strong diversity motivations, a different dynamic emerges in which a single mega-platform integrates both opinion groups while smaller platforms persist as marginalized echo chambers. These results illustrate how platform-level structures arise endogenously from individual learning dynamics and highlight the conditions under which digital publics fragment or remain integrated.

Our results echo a key insight from classical segregation models (Sakoda, 1949; Schelling, 1969, 1971), which showed how individual motivations can aggregate into unintended macro-level outcomes. While Schelling demonstrated how even mild preferences for like-minded neighbors produced sharp spatial segregation, similar insights were already developed in Sakoda's pioneering work (Sakoda, 1949), recently revisited in Hegselmann (2017). In a similar spirit, our model shows how motivations for social approval or diversity can aggregate into unintended collective outcomes such as polarized online spaces. Unlike in classical segregation models, these outcomes constitute a social dilemma: individually rational platform choices can lock the system into collectively suboptimal states in which users fail to satisfy their preference for diversity and are globally worse off. This social dilemma, arising at the level of platform composition, emerges in our model because it

goes beyond fixed interaction structures and captures the co-evolution of user preferences and platform ecologies (see [Figure 1](#)). In this respect, our findings connect the segregation tradition with the literature on social dilemmas, where feedback-driven learning has long been recognized as a mechanism through which individually rational behavior can generate collectively suboptimal equilibria ([Flache and Macy, 2002](#); [Macy, 1989](#)).

This paper makes several key contributions to the study of platform dynamics and user behavior on social media. First, whereas most previous ABMs of online polarization focus on opinion change through interaction ([Keijzer et al., 2018](#); [Keijzer and Mäs, 2022](#); [Kozitsin, 2022](#); [Maes and Bischofberger, 2015](#)), we shift the analytical lens to how users adapt their communicative behavior within dynamic online environments. This perspective allows us to examine how a polarized and fragmented online public ([Dahlberg, 2007](#); [Sunstein, 2018](#)) can emerge even in the absence of belief change. By treating opinions as relatively stable dispositions and focusing instead on participation as an adaptive response to social feedback, our approach helps close the gap between theoretical opinion dynamics models and empirical work in computational social science ([Flache et al., 2017](#); [Sobkowicz, 2009](#)), which largely relies on digital trace data and emphasizes participation patterns across ideological camps ([Arvidsson et al., 2024](#); [Theocharis and Jungherr, 2021](#)). By simulating the evolving interplay between user preferences and platform features, the model provides a novel framework for studying how online segregation emerges from individual communication decisions, thereby preparing the ground for bridging abstract modeling traditions with empirical research on online publics.

Second, beyond simulation, we conduct a rigorous mathematical analysis of the model using bifurcation methods. This enables us to formally characterize its equilibrium properties and identify critical transitions in system behavior, showing how variations in the balance between social approval and diversity-seeking (μ) lead to different, partly co-existing collective outcomes: from polarized echo chambers to a dominant integrated mega-platform. In this way, we highlight the tipping points at which polarization persists despite diversity-seeking motivations and where more inclusive environments can emerge. Moreover, by comparing the analytical predictions with systematic computational experiments, we establish the validity of the analytical model, demonstrating that it not only captures the qualitative dynamics but also provides accurate quantitative predictions across a wide range of parameters. This provides a theoretical basis for understanding platform segregation and diversity dynamics.

Third, the model is highly flexible and can accommodate a wide range of attributes, beyond political or ideological stances. In our framework, both

platforms and opinions are modeled at a high level of abstraction. Platforms are not tied to specific companies but represent structured online environments that differ in normative climate and user composition. This can refer to distinct social media sites as well as to communities or channels within a single site. For instance, the model can be applied to posting decisions within platforms like Reddit, where users choose between different communities (called “subreddits”) based on feedback such as upvotes or comments. Likewise, opinions are treated as abstract dispositions that may stand for political orientations, cultural values, or communicative preferences such as attitudes toward toxicity. While in this paper we focus on a simple version with binary opinions and homogeneous motivations, this flexibility makes the modeling framework a versatile tool for fine-tuning and calibrating models to specific online settings and for exploring how platform design and user incentives interact across different contexts.

Finally, our model serves as a testbed for exploring intervention strategies. Rather than treating polarization as an inevitable outcome, we examine how small and strategically targeted changes in the online environment can alter system dynamics. As a use case, we show that rewarding minority expression—akin to a minority game incentive (Challet et al., 2004)—can help a newly introduced platform disrupt segregation and stabilize integration. This perspective bridges theoretical modeling with practical considerations of platform governance and highlights how even modest interventions can open pathways out of the segregation trap.

The paper is organized as follows. We will first introduce the ABM and provide intuition on the model behavior by looking at paradigmatic simulation runs. We then derive a minimal model and analytically characterize the model transitions by bifurcation analysis. Section “Equilibrium selection” shows that these theoretical results fit with the full ABM and uses the model for intervention analysis. Finally, we discuss limitations and implications of our approach.

Agent-based model

Agent-based models (ABMs) provide a mechanism-based perspective on collective social dynamics by simulating how populations of artificial agents evolve in repeated interaction. While social feedback theory (SFT) describes collective dynamics as repeated communication games allowing in principle for game theoretic and dynamical systems analysis (Gaisbauer et al., 2020), ABMs provide full flexibility in devising such communication games. Here, the ABM serves as a benchmark that implements the full model logic, allowing us to study the resulting dynamics before introducing analytical

approximations. While we later show that a reduced dynamical system captures the core dynamics well, the ABM provides an essential foundation for validating this theoretical description.

In the model, we assume that a number of N users (called agents) decide to engage in opinion expression activity on M different social media platforms. Different platforms are indexed as $k \in \{1, 2, \dots, M\}$. Agents can also “opt out” by selecting $k = 0$, which represents non-participation in the current round. Formally, this corresponds to an empty environment that yields zero reward. We further assume that agents hold two opposing opinions ($o_i \in \{-1, 1\}$) on a single issue. These opinions are assigned at random at initialization and remain fixed throughout the simulation, dividing the population into N_+ agents holding opinion 1 and N_- agents with opinion -1 . This reflects a conceptual distinction between attitudinal dispositions and behavioral strategies. We may hence also think of the o_i 's as a stable cultural attribute or identity in the spirit of Schelling (1971) or Törnberg et al. (2021).

In subsequent rounds, agents choose to express their opinion on their currently preferred platform and evaluate that platform based on the social feedback obtained within it (rewards). In the next paragraphs, we discuss the model steps one by one. The full model code is provided in Appendix A.

Platform choice

Following Banisch and Olbrich (2019), our model uses reinforcement learning (RL) to model adaptive agent behavior. Specifically, in this model, preferences for platforms are represented via subjective expected utilities (i.e., Q -values) associated to the different options among which agents can choose. These evaluations are updated round by round based on past experiences (i.e., rewards). Let $Q_i(k)$ denote the Q -value agent i associates with platform k . Agents (indexed by i) choose platform k with a probability derived from the Boltzmann soft-max procedure, a method commonly used in decision-making models where options with higher expected rewards are chosen with higher probabilities. This approach balances exploration and exploitation, allowing agents to occasionally choose less-optimal platforms to explore new possibilities. The probability of choosing platform k is given by:

$$p_k^i = \frac{e^{\beta Q_i(k)}}{\sum_{l=0}^M e^{\beta Q_i(l)}}. \quad (1)$$

The parameter β , known as the exploration rate, governs the balance between randomness and rationality in decision-making. A value of $\beta =$

0 corresponds to purely random choices, while $\beta \rightarrow \infty$ implies fully deterministic choice, where agents always select the platform with the highest expected reward. Platform choice c_i is governed by equation (1) such that $c_i = k$ with probability p_k^i .

Recall that the first “environment”, $k = 0$, represents the option of non-participation in a given round. As we assign zero rewards to this action, $Q_i(0) = 0$ for all agents throughout the entire process. Note however that $Q_i(0) = 0$ does not imply that this option is never chosen (i.e. $p_k^i > 0$) as the choice probabilities in (1) are determined by the relative values of all $Q_i(k)$. In early stages or under high exploration (i.e., low β), non-participation may be chosen with non-negligible probability, especially when agents have not yet formed strong preferences for any specific platform.

Platform activity and opinion

After all agents made their choices, we compute the activity A_k and opinion O_k of each platform. The activity of platform k is defined as the number of users expressing their opinion on platform k in a given round

$$A_k = \sum_{i=1}^N \delta(c_i = k) \quad (2)$$

where $\delta(c_i = k)$ is an indicator function that equals 1 if agent i chose platform k , and 0 otherwise.

The overall opinion O_k on platform k represents the average opinion expressed by users on that platform. It is computed as the mean opinion of all active users on the platform in a given round:

$$O_k = \frac{1}{A_k} \sum_{i=1}^N \delta(c_i = k) o_i \quad (3)$$

where o_i represents the opinion of agent i . Platform activity and opinion are the most central variables for analyzing the temporal evolution of the ABM.

Rewards

In the model, users have two distinct motivations that make platform engagement appealing: (a) social approval of their expressed opinions and (b) exposure to opinion diversity on the platform. These incentives reflect the dual motivations that align with broader public sphere functions, where users may either seek affirmation within like-minded groups or pursue diverse

perspectives. We represent these motivations through two types of rewards. The first reward (a) assumes that agents perceive social support for their opinion as positive reinforcement. The average opinion on the chosen platform O_k reflects the likelihood of an agent receiving affirmative responses, capturing this social feedback as follows:

$$r_a^j(c_i = k) = O_k o_i. \quad (4)$$

Accordingly, when the average platform opinion O_k is negative, an agent with a negative opinion will receive positive feedback, while an agent with $o_i = +1$ will experience negative feedback.

But agents are motivated not only by support and the affirmation of seeing their opinions echoed. A platform might also lose appeal if there is insufficient opinion diversity, as a lack of contrasting views reduces the potential for argumentation and meaningful discussion. This second type of reward (b) represents the value agents place on encountering diverse perspectives, and is captured by

$$r_d^j(c_i = k) = 1 - |O_k|, \quad (5)$$

which decreases linearly with the opinionatedness of the platform in the current round. This diversity reward reaches a maximum of one when the platform consists of a balanced mix of agents from both opinion groups and drops to zero if all agents share the same opinion. In this way, r_d reflects the reward agents gain from a diverse online environment in terms of an increasing reward by exposure to the opinion of others.

We combine these two motivations introducing the basic model parameter μ by

$$r_k^j = (1 - \mu)r_a^j + \mu r_d^j. \quad (6)$$

The parameter μ models the tradeoff between seeking social approval and opinion diversity as a convex combination of the two reward contributions r_a and r_d . When $\mu = 0$, agents care exclusively about social approval (homophily), while when $\mu = 1$, they prioritize engaging with diverse perspectives. Intermediate values of μ capture a balanced social preference for both approval and diversity.

Agent update

The basic idea of social feedback models is that agents learn from the rewards obtained in one round and adjust their Q-values for the next round

accordingly. This is implemented via Q-learning, where the value $Q_i(k)$ associated with platform k is updated at the end of round t as follows:

$$Q_i^{t+1}(k) = Q_i^t(k) + \alpha(r_k^i - Q_i^t(k)). \quad (7)$$

That is, after agent i chooses platform k at time t , she receives the corresponding experienced reward r_k^i , and uses this to update her expected reward $Q_i(k)$ for the next round $t + 1$. The quantity $(r_k^i - Q_i^t(k))$ is known as the temporal difference error. Note that agents update only the Q -value of the platform they selected in the current round and do not observe rewards (or opinions) on platforms they did not join. The learning rate α governs how strongly agents adapt to new feedback: low values imply slow adjustment and stronger inertia, while high values increase sensitivity to recent experiences.

Platform rewards

In addition to the platform activity A_k and opinion O_k (2) and (3), we also keep track of the rewards agents obtain on average from a given platform. In analogy to the previous observable, we define the platform reward as

$$R_k = \frac{1}{A_k} \sum_{i=1}^N \delta(c_i = k) r_k^i \quad (8)$$

as the average reward agents that engage with platform k receive in the current round. The platform rewards thus trace how satisfied users are on a certain platform and reflect the extent to which a platform market serves the needs of users.

Code availability

A compact version of the model is provided as an Appendix. The full simulation code and additional analysis scripts are available at: <https://github.com/sevenbanisch/A-dynamical-model-of-platform-choice>.

Simulations

Model phenomenology

Simulation setting. Round by round, all agents choose one out of M platforms based on an evaluation $Q_i(k)$. For the simulation within this section, we use

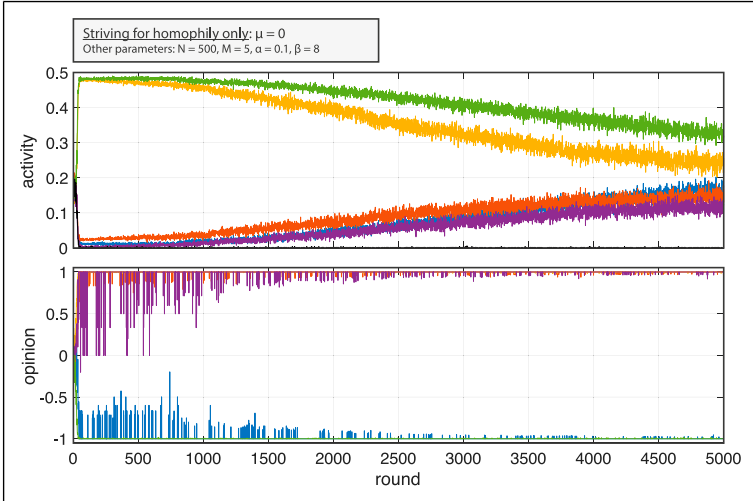


Figure 2. Example run for $\mu = 0.0$ which shows platform polarization due to homophily, with opinions clustering on separate platforms. The evolution of the normalized platform activity for the five platforms ($A_k(t)/N$) is shown in the top panel, and the platform opinion ($O_k(t)$) is shown below. The share of non-participating agents ($k = 0$) is also shown in the activity plot by the black curve (very close to zero). Other parameters: $N = 500$, $M = 5$, $\alpha = 0.1$, $\beta = 8$.

$N = 500$ agents that can choose among $M = 5$ different platforms with parameters $\alpha = 0.1$ and $\beta = 8$. At the start of the simulation, all Q -values are initialized to zero, meaning that all platforms (including the non-participation option $k = 0$) are chosen with equal probability. Opinions $o_i \in \{-1, 1\}$ are assigned at random so that, on average, half the population supports one and the other half the other opinion ($N_+ \approx 250$, $N_- \approx 250$). Our variable of interest is the preference for diversity μ .

Preference for homophily only ($\mu = 0$). Let us first consider the extreme cases, starting with agents who derive reward from seeing their opinion supported on the platform and do not strive for diversity. This corresponds to $\mu = 0$ in equation (6). In Figure 2 we show an exemplary run for this case. The evolution of the normalized platform activity for the five platforms ($A_k(t)/N$) is shown in the top panel, and the platform opinion $O_k(t)$ is shown below. The share of non-participating agents ($k = 0$) is also shown in the activity plot by the black curve and remains close to zero, as echo chambers provide maximum reward (≈ 1) at $\mu = 0$.

In this setting, the system quickly approaches a state where agents with a positive opinion meet on three platforms (yellow, orange and violet) and agents with a negative opinion gather on the remaining two (green and blue). We shall refer to this outcome as platform segregation, as agents with opposing opinions have a very low chance of interacting on the same platform. This separation of communicative environments mirrors key aspects of polarization, particularly the breakdown of cross-cutting exposure, even though the underlying opinions themselves remain fixed. Notice in the activity plot that two platforms initially attract a large user share so that almost all agents prefer one of the two. In particular, agents with a negative stance gather on the green platform and agents with $o_i = 1$ on the yellow one. The remaining platforms, though strongly opinionated from the beginning, are very low in activity, but they attract more and more users as time evolves. Eventually, in this run, platform activity will converge to a stationary profile in which agents with the same opinion choose the respective platforms on their side with equal probability ($\approx 1/4$ for green and blue, and $\approx 1/6$ for yellow, orange and violet). The convergence of Q-values causes agents to choose among platforms that share their opinion uniformly at random.

Notice further that we may observe model realizations in which only a single platform emerges on one side of the opinion spectrum, and the other four share users on the other side (not shown). The activity then approaches a state in which half on the users (i.e. those with the respective opinion) gather on one platform, whereas the other opinion group engages at equal share on the remaining ones ($\approx 1/8$).

Striving for diversity only ($\mu = 1$). The second limiting case is when agent only opt for diversity, and draw no reward from seeing their own opinion backed by others. That is, $\mu = 1$. An example run of this situation is shown in [Figure 3](#).

In this parameter setting, we observe that all platforms approach an average opinion $O_k \approx 0$ and are hence not opinionated. The activity approaches a value of $A_k/N \approx 1/5$ meaning that agents do not strongly prefer one platform over another. Again, we remark that the Q-values are approximately equal for all k options so that the platform choice probabilities are equal. The number of non-participating agents is negligible, as participation on balanced platforms provides maximum reward (≈ 1) for $\mu = 1$.

Striving for diversity and homophily ($\mu = 0.7$). A series of interesting phenomena may emerge in between these two limiting cases when mixing rewards from affirmative feedback (social approval) and diversity by the parameter $\mu \in [0, 1]$. We illustrate this by looking at a parameter value $\mu = 0.7$ by which agents




Figure 3. Example run for $\mu = 1.0$, illustrating balanced platform activity with diverse, non-opinionated platforms when agents prioritize diversity. In this parameter setting, all platforms approach an average opinion $O_k \approx 0$ and are hence not opinionated. Other parameters: $N = 500$, $M = 5$, $\alpha = 0.1$, $\beta = 8$.

draw a considerable reward from interaction with agents from the other group, but also value approval. In [Figures 4](#) and [5](#) we compare two model runs with $\mu = 0.7$. In this case, we show only the first 1000 rounds of the model, as it quickly settles into a stationary state.

The parameter $\mu = 0.7$ is chosen because the model can enter two very different equilibrium states. Even under the same initial assignment of Q-values (all $Q_k^0 = 0$), the stochastic process implemented by the model may either approach a state of platform polarization (as with $\mu = 0$), or a state in which one maximally diverse platform dominates over four marginal but opinionated platforms.

In the first realization (run A), platform polarization is observed after an initial phase in which two platforms compete to provide a diverse space. For up to 150 rounds, the orange and the blue platform remain in the middle of the opinion scale. This means that an approximately equal share of users from both opinion groups express their opinion on them. These platforms are also the most active during that phase. But, in this run, this situation becomes unstable so that one of the platforms is drawn to the positive and the other one to the negative side. After 400 rounds, no platform affords opinion diversity and the activity levels converge. In this run, two platforms (yellow and blue)




Figure 4. First run for $\mu = 0.7$, demonstrating polarization after an initial phase of platform competition (referred to as Run A). Other parameters: $N = 500$, $M = 5$, $\alpha = 0.1$, $\beta = 8$.

are populated by agents with a positive opinions, and three platforms emerge on the negative side. Consequently, their activity levels are around $1/4$ and $1/6$ respectively. However, notice that in this setting, the probability of non-participation (i.e., selecting $k = 0$) is clearly above zero as echo chambers provide a relatively low reward $\approx (1 - \mu) = 0.3$ which is only slightly higher than the value $Q_i(0) = 0$ for non-participation.

The second run (run B) shows a qualitatively different trend. Here, a single platform (blue) provides a space for interaction for agents with opposing opinions. After a short period in which the platform becomes more populated by agents with a positive opinion (rounds 100 to 300), the platform stabilizes as a space on which agents from both sides express opinions. With an activity of more than 80%, the blue platform grows very large and we shall refer to this as a mega-platform. The remaining platforms are all clearly on one side of the opinion scale connecting agents with equal opinions. But they remain marginal, as their activity converges to a low level.

Complex transient dynamics. The model exhibits rather complex transient phenomena which is shown by a realization with a slightly increased $\mu = 0.75$ in Figure 6. As before, this run features the emergence of a mega-

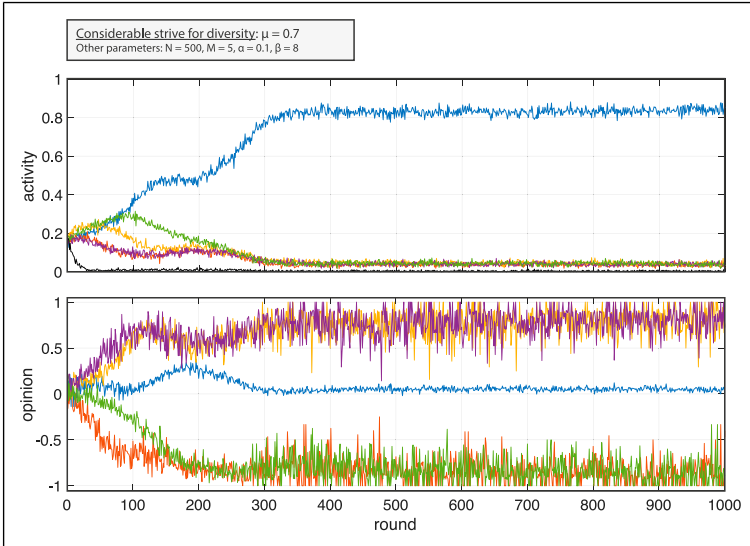


Figure 5. Second run for $\mu = 0.7$, showing the emergence of a single diverse platform (mega-platform) while others become marginalized (referred to as Run B). Other parameters: $N = 500$, $M = 5$, $\alpha = 0.1$, $\beta = 8$.

platform on which around 80% of activity takes place. However, the process leading to this final stationary situation is highly non-trivial. First, two highly active platform featuring diverse opinions emerge (green and violet in this case). One platform (orange) stabilizes on the negative and the remaining two (blue and yellow) on the positive side. Then both most active platforms tend to opposing opinion sides affording less diversity. Only the green platform survives this competition and stabilizes at a high level of over 80 % of activity.

Platform rewards

In this section, we return to the two model realizations with $\mu = 0.7$ and compare the two qualitatively different outcomes with respect to how rewarding the emergent platform market is for users. We base our definitions on the actual user experiences in the simulation, tracking the rewards agents obtained during in each round (equation (8)). For each single platform, the upper two panels in Figure 7 show the average reward agents experience when active on the respective platform.

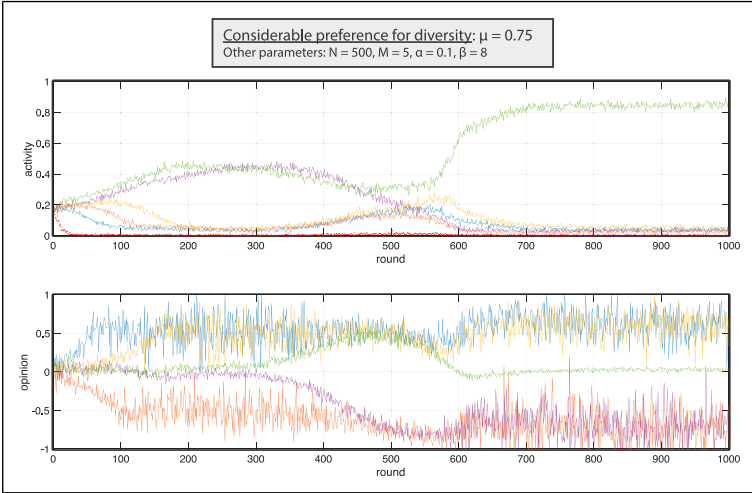


Figure 6. An example run which highlights complex transitions, with diversity initially emerging before one dominant platform forms ($\mu = 0.75$). Other parameters: $N = 500$, $M = 5$, $\alpha = 0.1$, $\beta = 8$.

Run A shows the system's tendency to segregate into polarized echo chambers, where agents only interact with like-minded peers and receive low rewards. Two initially diverse platforms collapse after an initial period, user rewards approach a value of $r_k = 0.3 = (1 - \mu)$ on all platforms. This corresponds to the reward that would be expected in a segregated and polarized platform market. Agents only encounter like-minded others and draw a reward of $(1 - \mu)$ from these interactions.

In contrast, Run B illustrates the emergence of a single dominant, diverse platform, where agents benefit from cross-opinion interactions, resulting in higher rewards. On the single mega-platform, the experienced reward is close to $r_k \approx 0.7 = \mu$. This value corresponds to the theoretically expected reward if all agents meet in perfectly balanced platforms. The remaining four platforms are marginal, but strongly opinionated so that opinion expression on them features rewards $r_k \approx 0.3$ of social approval.

This user-centric perspective on the degree to which emergent platform markets serve the needs of agents allows for a qualification of different platform constellations into optimal and suboptimal equilibria. Clearly, if $\mu = 0$ all agents strongly prefer interaction with agents of the same opinion. Hence, agents are maximally satisfied if they meet in homogeneous echo chambers that emerge under this parameter. More precisely, if $0 \leq \mu \leq 0.5$ in

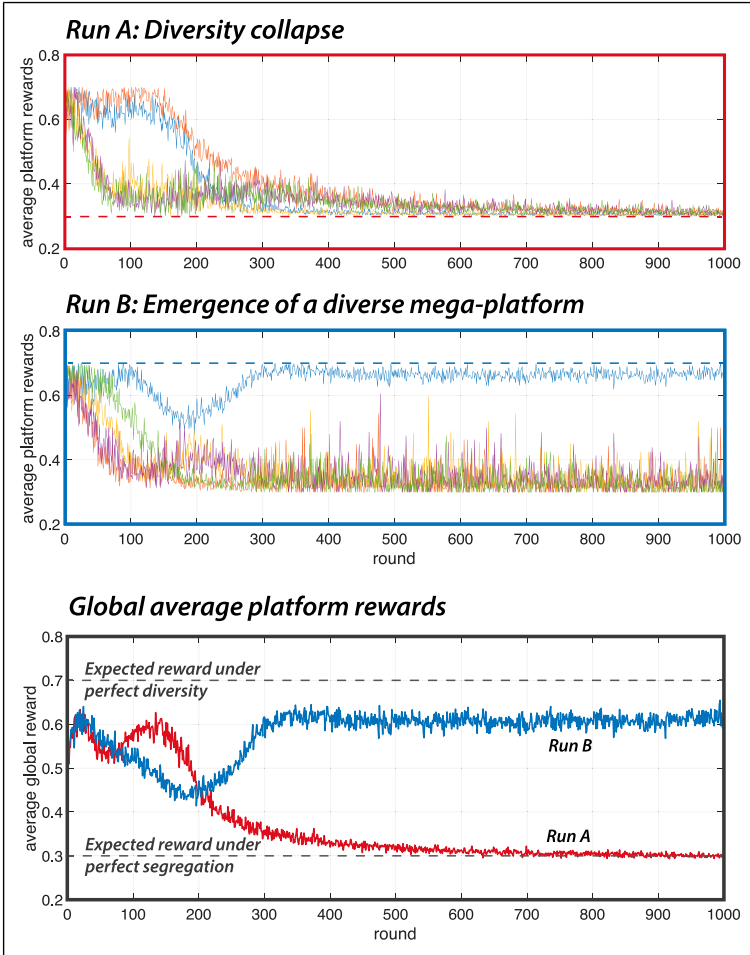


Figure 7. Comparing the two runs with $\mu = 0.7$ with respect to how rewarding the emergent platform market is for users. The dashed lines mark the expected rewards given $\mu = 0.7$ under the extreme (equilibrium) conditions of perfect diversity (global optimum) and complete segregation (suboptimal equilibrium).

equation (6), agents experience a polarized platform market maximally positive and can expect a reward $E_{\mu \leq 0.5}^*(r) = (1 - \mu)$ from engaging in opinion expression. For values $\mu > 0.5$, on the contrary, the optimal outcome would be a platform market of perfect diversity. The maximally expected reward if all agents meet in perfectly balanced online spaces is $E_{\mu > 0.5}^*(r) = \mu$.

Perfect diversity would be the optimal solution under the parameter value which lead to run A and B.

The global average reward (Figure 7, bottom panel) can therefore be interpreted as a user-centered measure of system-level performance. In game-theoretic terms, it functions as a social welfare metric, capturing how well the current distribution of users across platforms satisfies their individual preferences. As is clearly visible in Figure 7, when $\mu = 0.7$ agents are on average significantly more satisfied by a market with a single mega-platform that enables inter-opinion exchanges (run B) compared to a segregated situation (run A).

This illustrates a broader point developed analytically in the next section: under identical preferences, platform choice dynamics can settle in qualitatively different equilibria with varying levels of collective satisfaction.

Analytical characterization

The previous section has shown that the platform choice model exhibits complex dynamics and features a set of different final equilibrium states at different parameters values μ . In this section, we aim at a qualitative understanding of the global behavior of the model based on dynamical systems theory.

Model formulation

We first describe how to analytically formulate the model dynamics as a system of differential equations. For this purpose, we study the model with binary opinions ($o_i \in \{-1, 1\}$), as in in the previous section, but restrict to two competing platforms ($M = 2$). To simplify the analytical treatment, we exclude the option of non-participation ($k = 0$), reducing the number of available actions to two. We analyze the dynamics at a group level. That is, we assume homogeneity over the group of supporters ($\forall i$ with $o_i = 1$) and the opposing opinion group ($\forall i$ with $o_i = -1$).

4D System. The system dynamics is then described by four Q-values $Q(k|o)$ with platforms $k \in \{1, 2\}$ and opinions $o \in \{-1, 1\}$. $Q(1|1)$, for instance, denotes the value agents with $o = 1$ associate with platform 1. For further convenience we shall denote this as $Q_1^+ = Q(1|1)$. Likewise, $Q(1|-1) = Q_1^-$ is the Q-value the other opinion group with $o = -1$ assigns to platform 1. Correspondingly, we write for platform two Q_2^+ and Q_2^- .

In each round, the opinion groups choose one of the platforms by a soft max that compares the associated Q-values (Q_1^+ and Q_2^+ for agents with opinion +1). That is,

$$\begin{aligned} p_1^+ &= \frac{1}{1 + e^{\beta(Q_2^+ - Q_1^+)}} \\ p_2^+ &= \frac{1}{1 + e^{\beta(Q_1^+ - Q_2^+)}} \end{aligned} \quad (9)$$

for the group with $o = +1$, and equivalently for group $o = -1$ by replacing the superscript + with $-$. These expressions follow directly from equation (1) by applying the soft max choice rule to the case of two platforms ($M = 2$) and assuming group-level homogeneity.

The two platform opinions O_1 and O_2 can be computed on that basis by

$$O_k = \frac{p_k^+ - p_k^-}{p_k^+ + p_k^-}, \quad (10)$$

where $p_k^+ - p_k^-$ captures the relative prevalence of one opinion over the other and $p_k^+ + p_k^-$ is the platform activity A_k . The dynamics of the system are then defined by a 4D system of differential equations

$$\begin{aligned} \dot{Q}_1^+ &= \alpha p_1^+ (r_1^+ - Q_1^+) \\ \dot{Q}_2^+ &= \alpha p_2^+ (r_2^+ - Q_2^+) \\ \dot{Q}_1^- &= \alpha p_1^- (r_1^- - Q_1^-) \\ \dot{Q}_2^- &= \alpha p_2^- (r_2^- - Q_2^-), \end{aligned} \quad (11)$$

where α is the learning rate. The rewards are given as before by

$$\begin{aligned} r_k^+ &= (1 - \mu)O_k + \mu(1 - |O_k|) \\ r_k^- &= -(1 - \mu)O_k + \mu(1 - |O_k|) \end{aligned} \quad (12)$$

Note that in the full 4D formulation (11), the Q-value update is weighted by the probability p_k^\pm of selecting a given platform. This reflects that agents only update the Q-value of the platform they actually selected, i.e., learning is conditioned on observed outcomes. The resulting dynamics are thus frequency-dependent: the rate at which each Q-value is updated depends on how frequently the corresponding action is taken.

2D System. We can further reduce the system by looking at the differences of Q-values for each opinion group. For this purpose, we define

$$\begin{aligned}\Delta Q^+ &= Q_2^+ - Q_1^+ \\ \Delta Q^- &= Q_2^- - Q_1^-\end{aligned}\tag{13}$$

capturing to what extent both opinion groups respectively prefer platform 2 over platform 1. We exploit the fact that the selection probabilities (1) depend only on ΔQ . With that, the platform opinion O_k (10) and the expected rewards r_k^+, r_k^- (12) can be computed as before.

In order to close the dynamical equations in terms of ΔQ^+ and ΔQ^- , we have to make an assumption which differs from the 4D case and the ABM. Namely, we assume that agents always observe both rewards, regardless of which platform they choose. This corresponds to removing the frequency dependence (i.e., setting $p_k^\pm = 1$) in (11) and yields analytically tractable dynamics that still match the ABM very well when counterfactual learning is included.¹ With this, we obtain

$$\begin{aligned}\dot{Q}_2^+ - \dot{Q}_1^+ &= \Delta \dot{Q}^+ = \alpha(r_2^+ - r_1^+ - \Delta Q^+) \\ \dot{Q}_2^- - \dot{Q}_1^- &= \Delta \dot{Q}^- = \alpha(r_2^- - r_1^- - \Delta Q^-).\end{aligned}\tag{14}$$

The formulation of the model in form of a system of coupled differential equations provides a rather complete understanding of the model dynamics including the equilibrium states (fixed points) and the impact of initial conditions. A global understanding of the impact of μ can be obtained with bifurcation analysis, recovering the complex structure of critical transitions observed in the simulations.

Phase portraits and ABM comparison

We will first look at the phase portrait (or phase plot) of the 2D system for a given μ . A phase portrait shows for any combination $(\Delta Q^+, \Delta Q^-)$ in which direction the dynamical process will evolve under the model rules. Let us briefly recall how to interpret the two evolving variables $(\Delta Q^+, \Delta Q^-)$. If ΔQ^+ is positive, this means that positive opinion group ($o = 1$) favors platform 2 over 1, whereas a negative value indicates that platform 1 is favored. Likewise, the other opinion group ($o = -1$) prefers platform 2 if $\Delta Q^- > 0$. If $\Delta Q^+ < 0$ and $\Delta Q^- < 0$, for instance, both groups prefer posting on platform 1. The dynamical system (14) describes how these relative preferences evolve in time. In Figure 8, we show these temporal changes as a vector plot for $\mu = 0.7$, where simulations revealed the co-existence of two kinds of equilibria.




Figure 8. Phase portrait of the analytical model and comparison with 500 trajectories of an aligned ABM for $\mu = 0.7$ and $\beta = 8$. The blue and yellow curves show the null-clines of the two differential equation (14). The fixed points correspond to the intersections of these curves and the stable ones are marked by the circles (red for platform segregation, gray for a single dominant platform). In blue, 500 independently generated ABM trajectories are super-imposed. All model realizations are initialized at $(\Delta Q^+, \Delta Q^-) = (0, 0)$ and run for 1000 rounds. Note that all runs terminate near stable fixed points and the six different equilibria are visited with non-zero probability.

Note that phase plots show how the system would evolve for any specific initial condition. In a numerical simulation model, generating such a global picture would require a large number of repeated simulations with systematically varied initial conditions. Moreover, phase plots enable a graphical solution with respect to the different equilibrium states to which the model may settle. To identify these equilibria, we visualize in Figure 8 the null-clines with $\dot{\Delta Q}^+ = 0$ and $\dot{\Delta Q}^- = 0$ superimposed on the phase plot. Their intersections indicate the fixed points of the system, representing

configurations in which no further change is expected. These fixed points can be either stable or unstable; stability is assessed through the Jacobian matrix, as described in the [Online supplementary materials](#).

This analysis accurately recovers the complex structure of equilibria observed for $\mu = 0.7$ in the simulation section. There are six stable points for the analytical two-platform model. Four of them correspond to the mega-platform constellation where both opinion groups strongly prefer the same platform, marginalizing others to extremity. In our case of two platforms, the secondary platform is also strongly opinionated, because only one group interacts on it at a significant rate, while the other group avoids it altogether. This is reflected in the symmetry of the fixed points with respect to the diagonal. The remaining two stable fixed points, marked red, correspond to equilibria where one group prefers platform 1 and the other group platform 2 (i.e. $\Delta Q^+ < 0$ and $\Delta Q^- > 0$ or $\Delta Q^- < 0$ and $\Delta Q^+ > 0$).

To provide further justification that the analytical formulation of the models captures ABM dynamics, we also show 500 realizations of an ABM with two platforms in [Figure 8](#). The model has been aligned, omitting non-participation and including counter-factual learning as assumed when reducing from the 4D to the 2D version of the analytical model. This aligned ABM is hence completely consistent with the assumptions of the 2D analytical model and used for all simulation results reported in this section. All model runs are started with all Q-values at zero so that the two platforms are chosen with equal probability in the beginning. To project the simulations into the 2D phase plot, we measure the mean ΔQ_i^+ and ΔQ_i^- over the two opinion groups. We observed that all trajectories quickly leave this initial point at the center $(\Delta Q^+, \Delta Q^-) = (0, 0)$ and approach one of the stable fixed points. Interestingly, because of stochastic exploration in the ABM, the 500 simulation runs do not always converge to the same attractor. Instead, each of the six graphically identified fixed points is approached by multiple simulation trajectories, demonstrating that all equilibria are dynamically accessible.

The global picture

While phase plots show how the platform choice model behaves for a given μ , bifurcation analysis provides a global overview of the system transitions as the balance between social approval and diversity changes. Bifurcations mark the critical parameter values μ at which the global system dynamics undergoes qualitative changes. System transitions – sometimes called tipping points – are characterized by the emergence of new fixed points, a change of their stability or vanishing ones. We leverage the MatCont software package

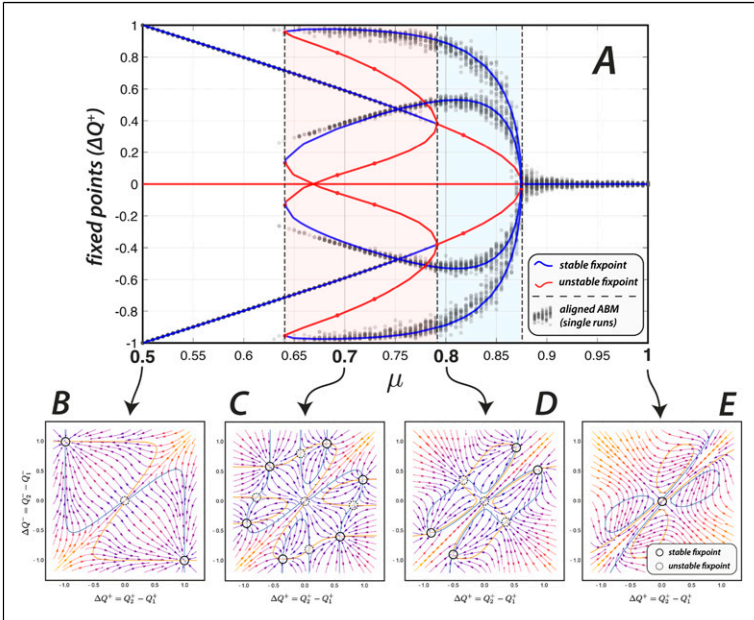


Figure 9. Global overview of the model behavior. Panel A: Bifurcation diagram tracing the fixed points as μ increases from 0.5 to 1. Blue corresponds to stable fixed points, red corresponds to unstable ones. Data from ABM simulations is shown by the gray dots. Each point shows the ΔQ 's in the stable state of a single model run (after $T = 1000$ steps, temporal average over the last 100 steps). For each $\mu \in [0.5, 0.505, \dots, 1.0]$, 100 model runs with $N = 1000$ agents are shown (10100 realizations). The matching between simulations and theory is remarkable. Panel B, C, D and E: Phase portrait of the system for four different values of μ . The blue and yellow curves show the null-clines of the two differential equation (14). The fixed points correspond to the intersections of these curves and are marked by the black circles (solid for stable, dashed for unstable fixed points).

(Dhooge et al., 2008) which works with a symbolic representation of the ODE system (14) and runs an associated ODE solver. The result is shown in the main Panel A of Figure 9 by the blue and red curves, corresponding to stable and unstable equilibria respectively.²

We again compare this analysis with a computational experiment performing systematic simulations of the aligned two-platform ABM. For each $\mu \in [0.5, 0.505, \dots, 1.0]$ (101 sample points), we run 100 simulations with $N = 1000$ agents. We measure the emergent stable state after $T = 1000$ rounds in terms of the mean ΔQ_i^+ and ΔQ_i^- defined over the two opinion groups

during the last 100 rounds. The final ΔQ 's for all 10100 model realizations are shown by the gray dots. We observe a remarkable matching between the ABM and the results of the bifurcation analysis.

Beyond the excellent fit with the aligned ABM, we find that the analytical model captures the series of non-trivial phenomena observed in the simulation of the full model as the alignment-diversity balance increases towards more diversity-seeking behavior. For this purpose, a series of phase portraits for specific values of μ are shown in Panels B to E of Figure 9. These values are chosen as representatives of qualitatively different model regimes as revealed by the bifurcation analysis.

Suboptimal platform segregation. In Panel B, the phase portrait for $\mu = 0.5$ is shown. There are three fixed points. The first one at $(0,0)$ is unstable so that the system is driven away from it towards one of the two stable fixed points at around $(-1, 1)$ and $(1, -1)$. In these states, the $+$ group strongly prefers one platform whereas the $-$ group prefers the other. In $(-1, 1)$, for instance, Q_1^+ is large compared to Q_2^+ so that the $+$ group interacts on platform 1 with a probability close to one. Conversely, $\Delta Q^- = 1$ indicates a strong preference of the $-$ group for platform 2. We are hence in a situation of segregation where platforms polarize into strongly opinionated echo chambers.

As shown in Section “Platform rewards”, this situation (equilibrium) is actually suboptimal as soon as $\mu > 0.5$. However, up to values of $\mu \approx 0.64$ these two equilibria are the only stable ones. This shows that the platform choice model will evolve into a suboptimal state of a segregated platform market despite the fact that all agents individually would be more satisfied by interaction on diverse online spaces in which they encounter opposing views.

The emergence of a second non-trivial attractor. As μ crosses a critical value of 0.64 (here exemplified by $\mu = 0.7$ in Panel C), the system bifurcates into a different qualitative regime and a series of new equilibria appears. There are four additional fixed points in which both opinion groups prefer the same platform to a slightly different degree. This nicely corresponds to the emergence of a mega-platform, on which agents with both opinions meet, as shown in the simulations section. There are four different such states in the two-platform version of the model. First, platform 1 or 2 may end up being chosen by both groups. Second, the remaining marginal platform may become populated predominantly by either the $+$ or the $-$ group. Hence, as in the simulations, this second platform is far less active, but strongly opinionated. The emergence of a stable state in which one diverse platform

dominates the others is one of the most surprising features of the platform choice model.

Coexistence of multiple equilibria. Still, the fixed points associated to platform polarization remain stable up to a value of $\mu \approx 0.79$. In this parameter regime, highlighted in Panel A by a light red shade, platform segregation and the emergence of a mega-platform coexist. We have observed this multi-stability in the aligned two-platform ABM as well as in the full model with several platforms. In this case (i.e. for $0.64 < \mu < 0.79$), the evolution of the model will strongly depend on initial conditions. Moreover, we might observe hysteresis: once the system enters a suboptimal attractor (platform segregation), it may be very difficult to return. On the other hand, the coexistence of these two equilibria provides ground for non-structural interventions such as promoting one platform as an space open to all groups, aiming to raise expectations (Q values) for one platform. We provide a systematic characterization of hysteresis effects and explore an intervention scenario in the next section using the full ABM.

Dominance as the only equilibrium. As μ crosses a value of 0.79 (exemplified by $\mu = 0.8$ in Panel D), the stable fixed points associated to platform segregation become unstable. The emergence of a single diverse mega-platform is the only stable outcome in this parameter regime (blue shaded region, $0.79 < \mu < 0.875$).

As far as we know, this state is unique to our model addressing segregation in a dynamic online environment. This outcome is surprising because one might expect that, under high diversity motivation, agents would evenly distribute across all platforms. Instead, a single integrated platform emerges as the dominant equilibrium, with other platforms marginalized to ideological extremes. The precise mechanism that leads to the survival of only a single diverse platform deserves further analysis. Notice, however, that a perfectly symmetric state where all platforms are equally integrated constitutes only a weak equilibrium: if users are indifferent between platforms, random fluctuations can lead to small asymmetries. Once such asymmetries arise, the reinforcement learning dynamics amplify the more rewarding (i.e., more integrated) platform, allowing it to dominate while the others decline. This self-reinforcing process may explain why only a single diverse platform survives.

Indifference. Finally, for $\mu > 0.875$ the fixed points associated with the mega-platform profile disappear and collapse into the fixed point at the center. $(0,0)$ becomes the only stable solution. This is exemplified in Panel E of [Figure 9](#)

for $\mu = 1$ which had also been considered in the simulation section as a limiting case. Q -values converge to the same values in the long run, meaning that agents are indifferent with respect to which platform they prefer. All platforms will be chosen with equal probability.

Equilibrium selection: Locking in and off the segregation trap

Social dilemmas typically involve the co-existence of multiple equilibria, only one of which is globally optimal (Macy and Flache, 2002). The bifurcation analysis revealed exactly such a structure: for $0.5 < \mu < 0.79$, a segregated collective outcome may emerge that no user prefers. This raises the question of when and how a system of learning agents—co-creating their own dynamic environment—approaches one equilibrium rather than another. Issues of path dependence and hysteresis are central here, since small differences in initial conditions or parameter changes can determine whether the system remains trapped in polarization or shifts toward integration.

The first purpose of this section is therefore to examine equilibrium selection. The mathematical reduction developed in the previous section provides valuable insights into the stability of alternative equilibria and the transitions between them. It shows how the system can lock in to segregation despite diversity-seeking motivations, and how it may eventually escape this trap when conditions change.

The second purpose is to validate the analytical results against the full multi-platform ABM. While the reduced two-dimensional model aligns well with the two-platform case, it remains an open question how far its predictions extend to the richer dynamics of the full system. In particular, we had to (i.) restrict to only two platforms and (ii.) assume that agents learn on the rewards of their non-choice. In Section “Simulations”, we already noted that multiple equilibrium paths exist at $\mu = 0.7$. Here we demonstrate that the mathematical reduction provides an accurate picture of the transition regime also for the full ABM.

To this end, we perform a series of computational experiments in two settings. First, we initialize the system in the “good” equilibrium with one dominant and diverse platform and study under what conditions it locks into segregation. Second, we initialize the system in the segregation equilibrium and examine when it locks off and converges toward the mega-platform attractor. We refer to these two settings as the lock-in and lock-off scenarios, respectively (see Figure 10). By varying μ and comparing the ABM outcomes with the predictions of the reduced analytical model, we assess both




Figure 10. Results of the lock-in (blue) and lock-off (red) experiments in comparison with the reduced analytical model using the PI as an order parameter. Dots indicate outcomes of the full ABM, while the bold dashed lines correspond to the PI evaluated at the equilibrium of the two-dimensional reduction. Both approaches reveal a coexistence regime ($0.64 < \mu < 0.79$) in which segregation and integration are simultaneously stable, giving rise to a hysteresis effect: once the system is in one equilibrium, it remains there until μ crosses a critical threshold.

the stability of the equilibria and the accuracy of the analytical approximation.

Computational experiment and model comparison

Activity-weighted variance. To compare the reduced two-platform model with the full multi-platform ABM, we measure segregation by the activity-weighted variance of platform opinions. For ease of reference, we call this the *Polarization Index* (PI). Let $A(k)$ denote the total activity on platform k and $O(k)$ the mean opinion expressed on that platform. With total activity $A = \sum_k A(k)$ and global mean opinion $\bar{O} = \frac{1}{A} \sum_k A(k) O(k)$, the PI is defined as

$$PI = \frac{1}{A} \sum_k A(k) (O(k) - \bar{O})^2.$$

The PI is low when both opinion groups are well represented on the same platform (integration) and high when activity is concentrated on separate platforms (segregation). It thus provides a consistent basis for comparing equilibria across models with different numbers of platforms.

Equilibrium initialization. To initialize the ABM directly in a desired equilibrium, we assign agents' Q -values to reflect the payoff structure of either the polarized or the integrated state. The construction of these initial Q -values

is guided by the patterns observed in simulation runs that converge to the respective equilibria. In the lock-off scenario, agents receive Q -values that favor participation on separate platforms aligned with their opinion, thereby producing multiple echo chambers. In the lock-in scenario, Q -values are chosen such that both opinion groups have an incentive to join the same dominant platform. To obtain a true equilibrium, we further impose preferences on the set of smaller platforms so that the second preference is shared among opinion groups (these platforms remain marginal but polarized). We apply an analogous procedure to initialize the analytical model near its corresponding fixed points. This ensures that ABM simulations begin in stable states of the analytical model, while still allowing us to test their robustness under variation of μ .

Computational setting. We sample ABM outcomes in a setting comparable to the two-platform experiments of the previous section. For each $\mu \in [0.5, 0.51, \dots, 1.0]$ (51 sample points), we run 100 simulations with $N = 1000$ agents and $M = 5$ platforms. We measure the system state after $T = 4000$ rounds in terms of the mean PI during the last 400 rounds. As in [Figure 9](#), we show the PI of all individual runs (corresponding to 2×5100 realizations for lock-in and lock-off initial conditions) as dots colored according to the initial condition. For the analytical model, we perform $T = 1000$ steps and compute the PI at the emerging fixed point. This system is deterministic and our setup basically traces the fixed points until they lose stability.

Hysteresis in the segregation dilemma

[Figure 10](#) reports the results of the lock-in and lock-off experiments in comparison with the analytical prediction. In the reduced model, both polarized and integrated equilibria are stable within the coexistence regime ($0.64 < \mu < 0.79$). The full ABM exhibits a similar hysteresis effect: the segregated equilibrium (red dots) persists up to $\mu \approx 0.76$ before losing stability, while the mega-platform equilibrium (blue dots) remains stable in the coexistence regime but disappears once $\mu < 0.6$. Both models therefore display the same qualitative transition sequence as μ increases: from segregation only, to coexistence, to dominance of a single integrated platform.

Beyond this qualitative match, the quantitative agreement between the reduced model and the full ABM is striking. The bifurcation thresholds predicted analytically closely align with the critical values observed in simulation, indicating that the simplified two-dimensional representation captures the essential dynamics of the full agent-based system. This path

dependence highlights the social-dilemma character of the model: although users may individually value diversity, collective dynamics can lock the system into persistent segregation even when more inclusive equilibria are attainable.

Sensitivity analysis: The role of α , β , and M

We finally assess the impact of other model parameters using the lock-off experiment for comparisons. The computational setting is as reported above, running 100 simulation trails for $\mu \in [0.5, 0.51, \dots, 1.0]$ (51 sample points) with $N = 1000$ agents. We explore the impact of α , β , and M by comparing the model results in terms of the final PI averaged over the 100 trails.

Learning rate and number of platforms. We first consider the influence of the learning rate α and the number of available platforms M . Since neither parameter enters the analytical reduction, the bifurcation structure derived in the previous section remains unaffected. [Figure 11](#) reports lock-off experiments for $\alpha \in \{0.1, 0.5\}$ and $M \in \{5, 8, 11\}$. The results confirm that the qualitative dynamics of equilibrium selection are robust: segregation persists in the coexistence regime and the system transitions to integration once μ exceeds the critical threshold, independent of learning rate or platform multiplicity.

Exploration rate. The exploration rate β controls how strictly agents follow their current best option, converging to pure best response as $\beta \rightarrow \infty$. In the opposite case of $\beta = 0$, all options are equally likely and the system evolves randomly. Unlike α and M , the parameter β enters directly into the analytical formulation, making it possible to compare simulation results with predictions of the reduced model.

[Figure 12](#) reports lock-off experiments (see [Figure 10](#)) for $\beta \in \{4, 6, 8, 10, 12\}$. As β decreases, the bifurcation threshold shifts to lower values of μ , and for $\beta = 4$ it almost vanishes. If values decrease further individual choices are random and the system dynamics is governed by these random fluctuations.

Accordingly, the misalignment between the ABM and the theoretical predictions increases at lower values of β . In this regime, platform choices are governed more by chance than by learned preferences, which amplifies stochastic variation and prevents the system from settling into the clear patterns predicted by the deterministic DE model.

Despite these stochastic fluctuations, the close correspondence between the bifurcation thresholds predicted analytically and the critical values observed in simulation further underscores the accuracy of the analytical




Figure 11. Results of the lock-off test for different learning rates $\alpha \in \{0.1, 0.5\}$ and number of platforms $M \in \{5, 8, 11\}$. We report the PI averaged over 100 simulation runs per μ . In all variations, the segregation equilibrium loses stability at values close to the theoretical prediction. The inset shows the same results focusing on the narrow transition region $0.72 < \mu < 0.82$. While there is no trend concerning the impact of M , the higher learning rate shifts the transition slightly to lower values of μ .

model. Together with the previous analyses of α and M , this demonstrates that the reduced model not only captures the qualitative dynamics but also provides accurate quantitative predictions across a wide range of parameter values, thereby establishing its theoretical validity.

Intervention experiments: Opening a new platform

The lock-off test provides a natural setting to study interventions, as it starts from a segregated equilibrium and asks under which conditions integration can emerge. Here we interpret the introduction of a new platform with neutral Q -values as an ecological intervention, opening a market niche by offering users an additional communicative environment not aligned with existing opinion camps. This alone already shifts the critical threshold of μ for successful integration by about 0.1 to the left (see orange curve in Figure 13), indicating that ecological changes in the platform environment can make diverse outcomes more likely even without changing user preferences. In the




Figure 12. Impact of the exploration rate on the lock-off test for $\beta \in \{4, 6, 8, 10, 12\}$. Analytical predictions (lines) are compared to ABM outcomes (dots) for $N = 1000$, $M = 5$, $\alpha = 0.1$, and $T = 4000$. For the ABM, the PI is averaged over 100 runs per μ (51 sample points), while for the analytical model we compute the PI of the respective fixed point.

following, we extend this idea by exploring protocols in which users of the new platform receive small additional rewards, allowing us to examine how modest design incentives can further increase the chances of escaping the segregation trap.

Modeling the emergence of a new platform. We implement this intervention in an ABM with $M = 5$ platforms, assigning segregation Q -values as in the lock-off scenario to four platforms and setting $Q = 0$ for the fifth. The same default parameters are used as in the lock-off test ($\alpha = 0.1$, $\beta = 8$, $N = 1000$) to allow direct comparison. In this setup, the system is initialized in a non-equilibrium low-reward state, reflecting the emergence of a new platform in an otherwise segregated environment.

Reward protocols. To explore the effect of design incentives, we consider six reward protocols that differ in scope, timing, and conditions.

- *Uniform incentives:*
 - (P1) all agents participating on the new platform receive a reward boost of 0.05.
- *Targeted incentives:*
 - (P2) only agents in the minority on the new platform receive the boost.

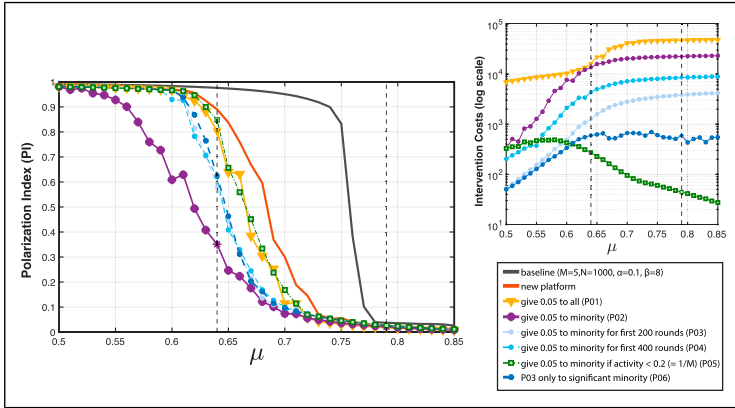


Figure 13. Intervention experiments in the lock-off setting with the introduction of a new platform ($M = 5$, $\alpha = 0.1$, $\beta = 8$, $N = 1000$). The figure compares six reward protocols: (P1) uniform boost for all users of the new platform, (P2) targeted boost for minority users, (P3–P4) temporary versions of P2 (200/400 rounds), (P5) conditional boost when activity on the new platform is below 20%, and (P6) conditional boost for genuine minorities. The figure reports the polarization index (PI, left) and the intervention costs (log scale, right) as a function of μ .

- *Temporary incentives:*
 - (P3) the targeted boost is applied only during the first 200 rounds.
 - (P4) the targeted boost is applied only during the first 400 rounds.
- *Conditional incentives:*
 - (P5) the targeted boost is given only if platform activity remains below 20% ($=1/M$).
 - (P6) the targeted boost is given only if the minority is significant ($|O(k)| > 0.05$).

This taxonomy highlights the flexibility of the ABM framework, which allows us to examine uniform, targeted, time-limited, and state-dependent interventions within the same modeling environment.

Intervention logic as platform strategy. The six reward protocols can be interpreted from the perspective of a platform founder deciding how to allocate a limited “kick-off budget”. For comparing the six protocols, Figure 13 considers how they shift the transition value at which the new platform survives as an integrated space. A naive strategy (P1) is to reward all users equally, which has only a marginal effect on integration (shift of about -0.02 in the lock-off test). Inspired by the logic of the Minority Game

(Challet et al., 2004; Challet and Zhang, 1997), a more efficient strategy (P2) is to reward only those who are currently underrepresented on the new platform. This is highly effective (shift of about -0.08), but also expensive. Restricting these minority incentives to the early phase (P3) is still effective (shift -0.05), although no structural change is observed once incentives stop after 200 rounds. Extending the horizon to 400 rounds (P4) brings no further benefit. Conditioning rewards on overall platform activity (P5) is largely ineffective (shift -0.02), suggesting that seeding participation early is more important than sustaining it later. Finally, targeting incentives only to clear minorities (P6) achieves the same effect as temporary incentives (shift -0.05) but at only a fraction of the cost, reducing expenditures by an order of magnitude. Taken together, these scenarios show how modest and strategically designed interventions can alter the equilibrium landscape, while highlighting trade-offs between effectiveness and cost.

Notice finally that only protocol P2 effectively alters the equilibrium structure of the collective communication game. Under P2, the integrated equilibrium becomes stable as soon as $\mu > 0.5$, whereas protocol P1 also modifies the structure but does not shift the critical threshold. In contrast, the temporal and conditional protocols (P3–P6) cease to operate once the desired state has emerged; the system then reverts to the original structure characterized analytically and numerically in previous sections. While P2 is costly if implemented as monetary rewards, its logic can also be realized through low-cost platform design features: for example, by framing participation as a “minority game” in which users receive symbolic rewards such as win messages or platform scores for daily contributions. This highlights that platform segregation is not inevitable but contingent on the availability of strategically designed interventions that reshape the equilibrium landscape of online publics.

Discussion

Our study builds upon Social Feedback Theory (SFT) which has been proposed as a framework for modeling how user behavior in online environments is shaped by social reinforcement (Banisch et al., 2022; Banisch and Olbrich, 2019). In repeated communication games agents express their opinions and receive feedback such as likes and downvotes from peers. While previous models developed within this paradigm (Banisch and Olbrich, 2019; Jacob and Banisch, 2023; Konovalova et al., 2023; Lefebvre et al., 2024) have focused on how users adapt their opinions based on the social feedback they receive—suggesting that social reinforcement may be the primary mechanism responsible for increasing polarization

(Lefebvre et al., 2024) and extreme opinion expression (Konovalova et al., 2023)—this work takes a different perspective. It focuses on the process of online community selection; an orthogonal but understudied facet of opinion dynamics which is especially relevant in the digital age.

In contrast to traditional opinion dynamics models, which primarily focus on how opinions evolve through direct interactions within fixed networks, our approach emphasizes the role of platform choice as a dynamic and strategic decision made by users. For instance, models like those discussed by Friedkin and Johnsen (2011), Flache et al. (2017) and Lorenz et al. (2021) typically assume that opinions change as users interact with others in pre-defined social networks, where exposure to like-minded or opposing views directly influences opinion shifts. These models often focus on opinion convergence or polarization within a single platform or network, without considering the broader context of multiple competing platforms. Our model departs from these frameworks by shifting the focus from opinion change to the choice of where to express opinions. This distinction is particularly relevant in an online context, where users have the freedom to navigate between platforms and adapt their behavior and online consumption based on the feedback they receive. By modeling how users with given opinions actively decide which platform to use based on rewards such as social approval or exposure to diverse views, we capture a key dynamic that is missing in many conventional opinion dynamics models addressing online phenomena (Keijzer et al., 2018; Keijzer and Mäs, 2022; Kozitsin, 2022; Maes and Bischofberger, 2015). This shift allows us to explore the interplay between user behavior and platform structures, offering new insights into how polarization and diversity can emerge across multiple platforms in an online ecosystem, rather than within a single, closed environment.

Our findings resonate with classical segregation models in the tradition of Sakoda and Schelling, which demonstrated how individual preferences can lead to unintended collective outcomes, such as spatial segregation (Sakoda, 1949; Schelling, 1969, 1971). In these models, even mild preferences for homophily result in sharp social divisions. Similarly, our model shows that even when users actively seek diverse interactions, online environments can still evolve into polarized, echo-chamber-like spaces. However, unlike classical segregation models, the segregated outcomes in our model constitute a genuine social dilemma, as individually adaptive behavior locks the system into states that are collectively inferior to feasible integrated alternatives. Moreover, the model also reveals an unexpected equilibrium where users become active on a single large platform, marginalizing other platforms into extreme opinion clusters. This state of equilibrium reflects how, under certain parameter values (such as a moderate μ), users gravitate toward a

platform that offers a sufficient mix of affirmation and diversity, potentially creating a “mega-platform” where the main discourse occurs, while alternative platforms become polarized and isolated. This co-existence of multiple qualitatively different equilibria, not found in earlier models, highlights the importance of the co-evolutionary dynamics between individual user behavior and macroscopic platform composition, adding new dimensions to the study of digital polarization. The presence of these equilibria also suggests a strong path dependence and potential for hysteresis, where initial conditions and past states significantly shape the system’s long-term outcomes.

While our model offers valuable insights into the dynamics of user behavior and platform choice, several limitations must be acknowledged. First, the model simplifies the complexity of user preferences by focusing primarily on two factors: social approval and exposure to diversity. In reality, users may be influenced by a broader range of factors, such as economic incentives, platform-specific features (e.g., ease of use, privacy concerns), or algorithmic transparency (Lorenz-Spreen et al., 2024). Studies have shown that users’ decisions to switch platforms may also be driven by factors like community norms or platform governance (Gillespie, 2018). In particular, the model currently centers on civic and deliberative motivations (Willaert and Olbrich, 2024)—such as seeking ideological diversity or social approval—but does not explicitly account for more interpersonal goals, such as maintaining friendships or sharing personal experiences, which also shape platform choice. These everyday motivations could be captured by modifying the reward structure within the social feedback framework. Moreover, recent work (Batzdorfer et al., 2026) demonstrates the importance of heterogeneous user motivations, showing that such engagement may reflect distinct behavioral profiles and levels of habituation. These findings highlight the need for models that accommodate diverse user types and reward structures. Future work incorporating a richer set of user motivations could explore how users driven by different communicative goals co-evolve within the same platform ecology, potentially leading to new forms of segregation or integration dynamics.

Likewise, our model does not explicitly represent within-platform social networks or interpersonal ties. As a result, it abstracts away from a major source of switching costs: the inability to carry over one’s social graph when migrating between platforms. This limitation is non-trivial, as users often face trade-offs between staying connected to their existing network and moving to platforms that better align with their communicative values or norms. More broadly, the model does not account for network externalities, where the utility of a platform increases with the number of active users—a

mechanism that can strongly reinforce dominant positions in the platform ecosystem. While such structural constraints are not captured in our framework, they play an important role in real-world platform choice and could be incorporated in future work by modeling personalized social networks or user-specific costs of disconnection.

At the same time, the model adopts an abstract notion of “platform”, allowing it to be fine-tuned to a range of online settings. In this view, platforms may refer to distinct providers or to subspaces within a single provider (e.g., subreddits on Reddit). This perspective enables the model to capture segregation dynamics that emerge through self-selection into different communicative environments. In this sense, it offers a stylized representation of compositional effects that may underlie within-platform polarization.

Furthermore, our model focuses on a homogeneous population of users divided along a single dimension (e.g., political opinion). In real-world online environments, users hold a diversity of opinions across multiple issues, which interact in complex ways (Baldassarri and Goldberg, 2014; Converse, 2006; DellaPosta, 2020). By modeling users as holding just one opinion or identity, we may overlook the potential for cross-cutting interactions across different issues, which can either exacerbate or mitigate polarization (Banisch and Olbrich, 2021; Baumann et al., 2021; Camargo, 2020). The computational framework is flexible enough to expand the model to account for users’ multi-dimensional identities—including political, cultural, and social preferences (Baldassarri and Bearman, 2007). In particular, it allows to directly incorporate heterogeneous user populations with opinions measured in surveys or on online social networking data (Gaumont et al., 2018; Peralta et al., 2024; Ramaciotti Morales et al., 2022). While complicating the computational analysis, designing empirical scenarios that account for more complexity at the level of users would allow for a more comprehensive understanding of how users engage with different platforms.

A central modeling choice in our approach is to treat opinions as fixed dispositions rather than dynamic states. This allows us to isolate the behavioral dimension of platform selection and examine how users gravitate towards different online environments based on social feedback. This assumption reflects a separation of timescales between belief formation and behavioral expression, and is empirically supported by studies showing that users’ opinions on key issues tend to remain stable over extended periods of online activity (Batzdorfer et al., 2025; Oswald et al., 2025). Against this background, we conceptualize opinions as stable preference structures that guide action, while the reinforcement learning process governs the adaptation of behavioral strategies. While many opinion dynamics models focus

on the transformation of beliefs through interaction (Flache et al., 2017; Friedkin and Johnsen, 2011; Lorenz et al., 2021), our model emphasizes how interaction patterns themselves emerge and stabilize around given opinion structures. We view this not as a limitation, but as a shift in analytical perspective. Importantly, it opens pathways for empirical validation: rather than relying on self-reported attitudes or opinion surveys, our framework aligns with recent approaches in computational social science that use behavioral trace data to study participation dynamics and ideological alignment across platforms or within online communities (Arvidsson et al., 2024; Theocharis and Jungherr, 2021). While our model remains stylized, its core mechanisms—adaptive platform selection based on social feedback—can be empirically explored in multi-platform ecosystems or within large platforms that offer multiple communities (e.g., Reddit). Future extensions could incorporate dynamic opinions and study the co-evolution of belief systems and platform ecologies, though such work would operate on a different conceptual, analytical, and empirical level.

Finally, while our model focuses on competition between multiple platforms, it does not fully capture the influence of platform algorithms beyond simple reward mechanisms. Real-world platforms employ complex algorithms that curate content in highly personalized ways (Bakshy et al., 2015). Incorporating more detailed models of algorithmic curation—such as a user feed—would allow for a deeper exploration of the role these algorithms play in shaping opinion dynamics and platform choice. This would also provide more actionable insights for platform designers seeking to mitigate polarization while maintaining user engagement (Bhadani et al., 2022; Lazer et al., 2018).

Nevertheless, the findings of our model have implications for both the design and management of digital platforms, as well as for broader societal concerns about online polarization and discourse. One of the key insights from our study is the coexistence of multiple equilibrium outcomes: platforms can evolve into polarized, echo-chamber-like environments or, alternatively, a single dominant platform can emerge where opposing opinions coexist, while smaller platforms become marginalized. This highlights the critical role of platform design and algorithmic moderation in shaping these outcomes. In practice, algorithms that prioritize user engagement—often by amplifying content that resonates strongly with existing user preferences—can contribute to the formation of polarized spaces. When users seek out platforms that align closely with their beliefs, social media ecosystems can naturally segregate into clusters of like-minded individuals, reinforcing the very dynamics of polarization that have been widely criticized (Bakshy et al., 2015; Sunstein, 2018). This suggests that algorithmic personalization,

though often deployed to enhance user satisfaction, may inadvertently fuel online segregation by drawing users into homophilous communities.

However, our model also points to opportunities for intervention. By adjusting platform features or modifying recommendation algorithms to promote exposure to diverse viewpoints, it may be possible to foster environments where a dominant, more inclusive platform prevails. For example, algorithms that intentionally balance between reinforcing social approval and offering diverse content could create spaces where users are more likely to encounter opposing views, potentially reducing polarization (Lorenz-Spreen et al., 2021; Willaert et al., 2021). However, while our model shows that exposure to diverse opinions can help reduce polarization within a dominant platform, it does not necessarily follow that smaller platforms can be prevented from becoming ideologically extreme through content moderation policies alone. In fact, smaller platforms may attract niche user groups precisely because of their more focused, often extreme, content, and the marginalization of these platforms into more radical spaces could be an inevitable consequence of market dynamics rather than content policies. Efforts to reduce the visibility of extreme content on larger platforms, for instance, might push certain user groups toward alternative platforms, further entrenching ideological divides online (Cinelli et al., 2021).

This highlights a significant challenge: while promoting content diversity and algorithmic balance on larger platforms could foster more inclusive spaces, it is unlikely that such efforts would fully prevent the concentration of extreme views on smaller platforms. As users seek spaces that align more closely with their ideological preferences, particularly when faced with moderation on larger platforms, smaller, less moderated platforms may become hubs for more polarized and extreme discourse. This suggests that platform designers and policymakers need to carefully consider the trade-offs involved in content moderation and algorithmic design, especially as they relate to the broader platform ecosystem.

Our intervention scenario demonstrates one such possibility: when a new platform enters a segregated ecology and strategically rewards minority expression, it can disrupt existing patterns of polarization and attract more diverse participation. This illustrates how subtle changes to feedback mechanisms can shift equilibrium outcomes without relying on content moderation, which often raises normative and political concerns. While costly if implemented monetarily, such interventions could also take symbolic or gamified forms, for instance through daily scores or visible recognition for users who engage in cross-cutting discourse. By altering the incentives that govern behavioral adaptation, such interventions offer a

promising route to counteract self-reinforcing segregation and foster integrative dynamics within competitive platform ecosystems.

Conclusion

In this paper, we introduced a dynamic model of platform choice, extending Social Feedback Theory to capture how users' decisions to engage with particular platforms contribute to the broader online ecosystem. Through a rigorous mathematical analysis, including bifurcation analysis, and agent-based simulations, we uncovered two key equilibrium outcomes: one where users polarize into echo chambers on separate platforms, and another where a dominant platform fosters diversity by accommodating opposing opinions, marginalizing all other platforms to extremity. Our model offers new insights into how user preferences for diversity or social approval drive the evolution of online environments.

This approach provides a new perspective on online polarization by focusing on the adaptive choices users with different opinions make across competing platforms, rather than restricting dynamics to a single platform environment. It highlights the complex interplay between adaptive user preferences and emergent platform structure, demonstrating that even when users value diversity, the feedback loops between platform selection and social feedback can still lead to segregation. In doing so, our model sheds light on the digital public sphere, illustrating how user motivations for either social reinforcement or diverse interactions can shape the kinds of discourse that unfold across platforms. Our findings underscore the co-evolutionary relationship between individual behavior and platform composition, revealing how user motivations, path dependencies, and initial conditions collectively influence the structure and inclusivity of online public spaces. In particular, we show how small shifts in incentive structures—such as rewarding minority participation—can alter systemic outcomes and support integrative configurations.

Overall, this study offers both theoretical and practical insights into the design of social media platforms. It suggests that platform choices driven by user preferences play a significant role in structuring digital public spaces, with implications for mitigating polarization. Rather than relying solely on content moderation or user belief change, our findings highlight the potential of behaviorally grounded interventions that modify the feedback dynamics guiding online choice. Future research will build on this framework by integrating more realistic user identities and algorithmic factors, contributing to the design of robust strategies for fostering a sustainable digital ecosystem.

Acknowledgements

This project has received funding from the European Union's Horizon Europe-programme under grant agreement No 101094752 (SoMe4Dem). We would like to acknowledge Prof. Naomi Leonard's research lab at Princeton University, whose support was instrumental in the mathematical analysis.

ORCID iDs

Sven Banisch  <https://orcid.org/0000-0002-1263-580X>

Dennis Jacob  <https://orcid.org/0009-0001-9086-8684>

Tom Willaert  <https://orcid.org/0000-0002-3879-6767>

Eckehard Olbrich  <https://orcid.org/0000-0002-3987-1271>

Ethical considerations

No ethical approval was required for this paper.

Funding

The authors disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This project has received funding from the European Union's Horizon Europe programme under grant agreement No 101094752 (SoMe4Dem). We would like to acknowledge Prof. Naomi Leonard's research lab at Princeton University, whose support was instrumental in the mathematical analysis.

Declaration of conflicting interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Supplemental material

Supplemental material for this article is available online.

Notes

1. This assumption does not have a qualitative impact on the behavior of the DE system and fits with an ABM of the two-platform setting. Moreover, if we include into the ABM that agents see the reward they would have obtained in the non-chosen environment (counter-factual learning), we obtain a quantitative match between analytical formulation and the simulation model.
2. Notice that we only visualize ΔQ^+ in the bifurcation plot because the system is symmetric and the plot for ΔQ^- would be equivalent. See [Online supplementary materials](#).

References

- Arvidsson M, Hedström P, Jarvis BF, et al. (2024) *On the Intersection of Analytical Sociology and Computational Social Science*. Handbook of Computational Social Science.
- Bakshy E, Messing S and Adamic LA (2015) Exposure to ideologically diverse news and opinion on Facebook. *Science* 348(6239): 1130–1132. <https://doi.org/10.1126/science.aaa1160>
- Baldassarri D and Bearman P (2007) Dynamics of political polarization. *American Sociological Review* 72(5): 784–811. <https://doi.org/10.1177/000312240707200507>
- Baldassarri D and Goldberg A (2014) Neither ideologues nor agnostics: alternative voters' belief system in an age of partisan politics. *American Journal of Sociology* 120(1): 45–95. <https://doi.org/10.1086/676042>
- Banisch S and Olbrich E (2019) Opinion polarization by learning from social feedback. *Journal of Mathematical Sociology* 43(2): 76–103. <https://doi.org/10.1080/0022250x.2018.1517761>
- Banisch S and Olbrich E (2021) An argument communication model of polarization and ideological alignment. *The Journal of Artificial Societies and Social Simulation* 24(1): 1. <https://doi.org/10.18564/jasss.4434>
- Banisch S, Gaisbauer F and Olbrich E (2022) Modeling spirals of silence and echo chambers by learning from the feedback of others. *Entropy* 24(10): 1484. <https://doi.org/10.3390/e24101484>
- Batzdorfer V, Banisch S and Oswald L (2025) Attitude shifts under ideal discussion conditions: stability, shift, and depolarization. *Poster presented at IC2S2*.
- Batzdorfer V, Samory M and Banisch S (2026) Why conspiracy theory communities endure: the interplay of feedback, habits, and community context. In: Proceedings of the 18th International Conference on Web and Social Media (ICWSM). Forthcoming.
- Baumann F, Lorenz-Spreen P, Sokolov IM, et al. (2021) Emergence of polarized ideological opinions in multidimensional topic spaces. *Physical Review X* 11(1): 011012. <https://doi.org/10.1103/physrevx.11.011012>
- Bhadani S, Yamaya S, Flammini A, et al. (2022) Political audience diversity and news reliability in algorithmic ranking. *Nature Human Behaviour* 6(4): 495–505. <https://doi.org/10.1038/s41562-021-01276-5>
- Boyd D (2014) *It's Complicated: The Social Lives of Networked Teens*. Yale University Press.
- Brady WJ, Jackson JC, Lindström B, et al. (2023) Algorithm-mediated social learning in online social networks. *Trends in Cognitive Sciences* 27(9): 765–776. <https://doi.org/10.1016/j.tics.2023.06.008>

- Camargo CQ (2020) New methods for the steady-state analysis of complex agent-based models. *Frontiers in Physics* 8: 103. <https://doi.org/10.3389/fphy.2020.00103>
- Challet D and Zhang Y-C (1997) Emergence of cooperation and organization in an evolutionary game. *Physica A: Statistical Mechanics and Its Applications* 246(3-4): 407–418. [https://doi.org/10.1016/s0378-4371\(97\)00419-6](https://doi.org/10.1016/s0378-4371(97)00419-6)
- Challet D, Marsili M and Zhang Y-C (2004) *Minority Games: Interacting Agents in Financial Markets*. OUP Oxford.
- Cinelli M, Morales GDF, Galeazzi A, et al. (2021) The echo chamber effect on social media. *Proceedings of the National Academy of Sciences* 118(9): e2023301118. <https://doi.org/10.1073/pnas.2023301118>
- Converse PE (2006) The nature of belief systems in mass publics (1964). *Critical Review* 18(1-3): 1–74. <https://doi.org/10.1080/08913810608443650>
- Dahlberg L (2007) Rethinking the fragmentation of the cyberpublic: from consensus to contestation. *New Media & Society* 9(5): 827–847. <https://doi.org/10.1177/1461444807081228>
- DellaPosta D (2020) Pluralistic collapse: the “oil spill” model of mass opinion polarization. *American Sociological Review* 85(3): 507–536. <https://doi.org/10.1177/0003122420922989>
- Dhooge A, Govaerts W, Kuznetsov YA, et al. (2008) New features of the software matcont for bifurcation analysis of dynamical systems. *Mathematical and Computer Modelling of Dynamical Systems* 14(2): 147–175. <https://doi.org/10.1080/13873950701742754>
- Flache A and Macy MW (2002) Stochastic collusion and the power law of learning: a general reinforcement learning model of cooperation. *Journal of Conflict Resolution* 46(5): 629–653. <https://doi.org/10.1177/002200202236167>
- Flache A, Mäs M, Feliciani T, et al. (2017) Models of social influence: towards the next frontiers. *Jasss-The journal of artificial societies and social simulation* 20(4): 2.
- Friedkin NE and Johnsen EC (2011) *Social Influence Network Theory: A Sociological Examination of Small Group Dynamics*. Cambridge University Press, Vol. 33. <https://doi.org/10.1017/cbo9780511976735>
- Fudenberg D and Levine DK (1998) *The Theory of Learning in Games*. MIT press, Vol. 2.
- Gaisbauer F, Olbrich E and Banisch S (2020) Dynamics of opinion expression. *Physical Review E* 102(4): 042303. <https://doi.org/10.1103/PhysRevE.102.042303>
- Garrett RK (2009) Echo chambers online? Politically motivated selective exposure among internet news users. *Journal of Computer-Mediated Communication* 14(2): 265–285. <https://doi.org/10.1111/j.1083-6101.2009.01440.x>

- Gaumont N, Panahi M and Chavalarias D (2018) Reconstruction of the socio-semantic dynamics of political activist Twitter networks-method and application to the 2017 French presidential election. *PLoS One* 13(9): e0201879. <https://doi.org/10.1371/journal.pone.0201879>
- Gillespie T (2018) *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions that Shape Social Media*. Yale University Press.
- Granovetter M (1978) Threshold models of collective behavior. *American Journal of Sociology* 83(6): 1420–1443. <https://doi.org/10.1086/226707>
- Hegselmann R (2017) Thomas c. schelling and james m. sakoda: the intellectual, technical, and social history of a model. *The Journal of Artificial Societies and Social Simulation* 20(3): 15. <https://doi.org/10.18564/jasss.3511>
- Homans GC (1974) *Social behavior: Its elementary forms*. Harcourt Brace Jovanovich.
- Izuma K, Saito DN and Sadato N (2010) Processing of the incentive for social approval in the ventral striatum during charitable donation. *Journal of Cognitive Neuroscience* 22(4): 621–631. <https://doi.org/10.1162/jocn.2009.21228>
- Jacob D and Banisch S (2023) Polarization in social media: a virtual worlds-based approach. *The Journal of Artificial Societies and Social Simulation* 26(3): 11. <https://doi.org/10.18564/jasss.5170>
- Keijzer MA and Mäs M (2022) The complex link between filter bubbles and opinion polarization. *Data Science* 5(2): 139–166. <https://doi.org/10.3233/ds-220054>
- Keijzer MA, Mäs M and Flache A (2018) Communication in online social networks fosters cultural isolation. *Complexity* 2018(1): 9502872. <https://doi.org/10.1155/2018/9502872>
- Knobloch-Westerwick S (2015) *Choice and Preference in Media Use: Advances in Selective Exposure Theory and Research*. Routledge.
- Konovalova E, Le Mens G and Schöll N (2023) Social media feedback and extreme opinion expression. *PLoS One* 18(11): e0293805. <https://doi.org/10.1371/journal.pone.0293805>
- Kozitsin IV (2022) Formal models of opinion formation and their application to real data: evidence from online social networks. *Journal of Mathematical Sociology* 46(2): 120–147. <https://doi.org/10.1080/0022250x.2020.1835894>
- Kuran T (1989) Sparks and prairie fires: a theory of unanticipated political revolution. *Public Choice* 61(1): 41–74. <https://doi.org/10.1007/bf00116762>
- Lazer D, Baum MA, Benkler Y, et al. (2018) The science of fake news. *Science* 359(6380): 1094–1096. <https://doi.org/10.1126/science.aao2998>
- Lefebvre G, Deroy O and Bahrami B (2024) The roots of polarization in the individual reward system. *Proceedings of the Royal Society B* 291(2017): 20232011. <https://doi.org/10.1098/rspb.2023.2011>

- Lindström B, Bellander M, Schultner DT, et al. (2021) A computational reward learning account of social media engagement. *Nature Communications* 12(1): 1311. <https://doi.org/10.1038/s41467-020-19607-x>
- Lopez-Pintado D and Watts DJ (2008) Social influence, binary decisions and collective dynamics. *Rationality and Society* 20(4): 399–443. <https://doi.org/10.1177/1043463108096787>
- Lorenz J, Neumann M and Schröder T (2021) Individual attitude change and societal dynamics: computational experiments with psychological theories. *Psychological Review* 128(4): 623–642. <https://doi.org/10.1037/rev0000291>
- Lorenz-Spreen P, Mønsted B, Hövel P, et al. (2021) Algorithmic versus social amplification of misinformation in online networks. *Proceedings of the National Academy of Sciences* 118(1): e2011999118.
- Lorenz-Spreen P, Arslan RC, Kozyreva A, et al. (2024) Real-time assessment of motives for sharing and creating content among highly active twitter users. *Preprint at PsyArXiv*. Available at: <https://doi.org/10.31234/osf.io/jf8w9>.
- Macy MW (1989) Walking out of social traps: a stochastic learning model for the prisoner's dilemma. *Rationality and Society* 1(2): 197–219. <https://doi.org/10.1177/1043463189001002003>
- Macy MW (1991) Chains of cooperation: threshold effects in collective action. *American Sociological Review* 56: 730–747. <https://doi.org/10.2307/2096252>
- Macy MW and Flache A (2002) Learning dynamics in social dilemmas. *Proceedings of the National Academy of Sciences* 99(suppl_3): 7229–7236. <https://doi.org/10.1073/pnas.092080099>
- Maes M and Bischofberger L (2015) *Will the Personalization of Online Social Networks Foster Opinion Polarization?* Available at SSRN 2553436.
- Matignon L, Laurent GJ and Le Fort-Piat N (2012) Independent reinforcement learners in cooperative markov games: a survey regarding coordination problems. *The Knowledge Engineering Review* 27(1): 1–31. <https://doi.org/10.1017/s0269888912000057>
- Oswald L, Schulz WS and Lorenz-Spreen P (2025) A collective field experiment disentangling participation in online political discussions. *SocArXiv*.
- Peralta AF, Ramaciotti P, Kertész J, et al. (2024) Multidimensional political polarization in online social networks. *Physical Review Research* 6(1): 013170. <https://doi.org/10.1103/physrevresearch.6.013170>
- Ramaciotti Morales P, Cointet J-P, Muñoz Zolotoochin G, et al. (2022) Inferring attitudinal spaces in social networks. *Social Network Analysis and Mining* 13(1): 14. <https://doi.org/10.1007/s13278-022-01013-4>
- Ruff CC and Fehr E (2014) The neurobiology of rewards and values in social decision making. *Nature Reviews Neuroscience* 15(8): 549–562. <https://doi.org/10.1038/nrn3776>

- Sakoda JM (1949) *Minidoka: An Analysis of Changing Patterns of Social Interaction*. Berkeley: University of California.
- Scharlach R and Hallinan B (2023) The value affordances of social media engagement features. *Journal of Computer-Mediated Communication* 28(6): zmad040. <https://doi.org/10.1093/jcmc/zmad040>
- Schelling TC (1969) Models of segregation. *The American Economic Review* 59(2): 488–493.
- Schelling TC (1971) Dynamic models of segregation. *Journal of Mathematical Sociology* 1(2): 143–186. <https://doi.org/10.1080/0022250x.1971.9989794>
- Shoham Y and Leyton-Brown K (2008) *Multiagent Systems: Algorithmic, game-theoretic, and Logical Foundations*. Cambridge University Press.
- Simon HA (1978) Rationality as process and as product of thought. *The American Economic Review* 68(2): 1–16.
- Sobkowicz P (2009) Modelling opinion formation with physics tools: call for closer link with reality. *The Journal of Artificial Societies and Social Simulation* 12(1): 11.
- Stroud NJ (2010) Polarization and partisan selective exposure. *Journal of Communication* 60(3): 556–576. <https://doi.org/10.1111/j.1460-2466.2010.01497.x>
- Sunstein C (2018) *# Republic: Divided Democracy in the Age of Social Media*. Princeton university press.
- Sutton RS and Barto AG (2018) *Reinforcement Learning: An Introduction*. MIT press.
- Theocharis Y and Jungherr A (2021) Computational social science and the study of political communication. *Political Communication* 38(1-2): 1–22. <https://doi.org/10.1080/10584609.2020.1833121>
- Thorndike EL (1911) *Animal Intelligence: Experimental Studies*. The Macmillan Company.
- Törnberg P, Andersson C, Lindgren K, et al. (2021) Modeling the emergence of affective polarization in the social media society. *PLoS One* 16(10): e0258259. <https://doi.org/10.1371/journal.pone.0258259>
- Tucker JA, Guess A, Barbera P, et al. (2018) Social media, political polarization, and political disinformation: a review of the scientific literature. *Political Polarization, Social Media, and Political Disinformation*.
- Van Raemdonck N and Pierson J (2022) Conceptueel kader voor wisselwerking van platformkenmerken, affordances en normen op sociale media. *Tijdschrift voor Communicatiewetenschap* 50(4): 358–383. <https://doi.org/10.5117/tcw2022.4.005.raem>
- Vollmer H (2013) What kind of game is everyday interaction? *Rationality and Society* 25(3): 370–404. <https://doi.org/10.1177/1043463113492307>
- Willaert T and Olbrich E (2024) Social media and the public sphere: an empirically-informed taxonomy of platform affordances. *SoMe4Dem Project Deliverable*.

<https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e510638760&appId=PPGMS>

Willaert T, Van Eecke P, Van Soest J, et al. (2021) An opinion facilitator for online news media. *Frontiers in Big Data* 4: 695667. <https://doi.org/10.3389/fdata.2021.695667>