Vol. 2, No. 3, pp.: 22-44, 2024

Mining Coevolutionary Patterns of Opinions and Behaviors in Multiplex Social Networks

Xiaoyu Wei[1], Zheng Lin[2], Hongbo Liu[3], Yifeng Zhao[4,*]

[1] School of Computer Science and Technology, Hangzhou Dianzi University, Hangzhou 310018, China

[2] School of Modern Posts, Nanjing University of Posts and Telecommunications, Nanjing 210003, China

[3] School of Mathematical Sciences, Anhui University, Hefei 230601, China

[4] School of Artificial Intelligence, Hebei University of Technology, Tianjin 300401, China

** yifeng.zhao@hebut.edu.cn*

Article Information	
Received	18 April 2024
Accepted	29 August 2024
DOI	https://doi.org/10.63646/datamind.2024.020303

Abstract

In real-world social systems, individuals form heterogeneous opinions on the same issue and continually adjust both their attitudes and overt behaviors through repeated interaction with peers. As online social platforms become deeply interwoven with offline life, the interplay between opinion diffusion and behavioral choice has grown more pronounced, and a single-layer perspective is no longer sufficient to capture the true trajectory of collective evolution. Motivated by this gap, this article studies the coevolution of opinions and behaviors over a multiplex (two-layer) social network and develops a unified data-mining framework for discovering coevolutionary patterns and stable regimes. In the baseline specification, the opinion layer follows a DeGroot-style weighted update rule, while the behavior layer integrates three influence channels operating in parallel: neighbor imitation, payoff-driven adaptation, and cognition-behavior consistency. Around this model, we first investigate the synchronous-update regime and characterize the joint evolution of the average opinion and the cooperator share, then study the asynchronous-update counterpart and compare the two regimes in terms of convergence speed, transient pathways, and steady-state outcomes. Three canonical topologies are examined: random graphs, small-world graphs, and scale-free graphs. Finally, the model is calibrated against a longitudinal household survey on pro-environmental attitudes and recycling behaviors. Results show that opinions and behaviors coevolve along strongly coupled trajectories; that the opinion-dependency parameter, the imitation weight, and the payoff sensitivity exert decisive influence on consensus formation and on the diffusion of cooperative action; that in most regimes the cooperator share stabilises before the average opinion does,

indicating a behavior-leads-cognition-lags pattern; and that scale-free and random networks tend to converge faster and to higher steady cooperator levels than small-world graphs. Synchronous and asynchronous mechanisms agree on long-run outcomes but differ markedly along the transient path. The simulated trends match the survey data within 4 percentage points, confirming that the proposed framework offers a practical and explanatory tool for mining coevolutionary patterns in real multiplex social systems.

Keywords: ***Multiplex social networks; opinion dynamics; behavioral coevolution; data mining; synchronous update; asynchronous update***

1. Introduction

Online social platforms have moved well beyond their original role as channels for casual exchange. They are now interwoven with employment, public health communication, environmental campaigns, political mobilisation, and consumer choice (Lazer et al., 2009; Watts, 2007). At the same time, the offline social fabric — workplaces, neighbourhoods, family circles — continues to shape what people believe and how they act. The empirical consequence is that the same individual is simultaneously embedded in several overlapping interaction structures, each carrying a distinct kind of social signal. An observation that drives the present study is that opinion formation and behavioural choice rarely move on the same timescale or through the same channel, yet they constantly reinforce or destabilise each other (Centola, 2010; Christakis and Fowler, 2007).

Single-layer treatments of social dynamics — purely opinion-based or purely behavioural — have produced an extensive analytical inventory, ranging from threshold cascade models and bounded-confidence rules to evolutionary-game frameworks (Granovetter, 1978; Hegselmann and Krause, 2002; Szabó and Fáth, 2007). However, they cannot easily address situations in which an individual's stated belief on an issue (for example, a favourable view of recycling) drifts in one direction while the corresponding overt action drifts in another. This decoupling is increasingly visible in environmental psychology, vaccination behaviour, public-health compliance, and political participation (Kollmuss and Agyeman, 2002; Stern, 2000). Multiplex network theory provides a natural mathematical scaffold for such phenomena: each layer encodes a distinct interaction modality, and inter-layer coupling encodes the cognitive link a person maintains between thinking something and doing it (Boccaletti et al., 2014; Kivelä et al., 2014; De Domenico et al., 2013).

This article makes three contributions. First, it formulates a coevolutionary framework on a two-layer network in which the opinion layer follows a DeGroot-style weighted-averaging rule and the behavior layer combines three explicitly identified mechanisms: neighbour imitation, payoff-driven adaptation, and a cognition-behavior consistency drive. The decomposition is designed for downstream pattern mining: each mechanism is tied to a distinct parameter that can be estimated, swept, or held fixed. Second, the article studies in parallel the synchronous and asynchronous update regimes and shows that they agree on stable states yet differ in the transient pathway, providing a controlled view of how update timing shapes observable patterns. Third, the framework is calibrated against a longitudinal household survey on pro-environmental attitudes and behaviors, which makes the data-mining loop concrete and falsifiable.

The remainder of the article is organised as follows. Section 2 surveys related work on opinion dynamics, evolutionary games on networks, and multiplex coevolution. Section 3 presents the modelling framework and introduces the notation. Section 4 develops the synchronous-update analysis, including parameter sweeps and convergence diagnostics on canonical topologies. Section 5 introduces the asynchronous-update counterpart and

compares the two regimes. Section 6 reports the empirical case on environmental behavior. Section 7 discusses implications for data mining on multiplex social systems, and Section 8 concludes.

2. Related Work

2.1 Opinion Dynamics on Networks

Modern opinion-dynamics research traces back to two foundational lines of work. The DeGroot (1974) consensus model treats each individual as a Bayesian-style averager who blends own beliefs with the weighted beliefs of neighbours, leading under mild connectivity conditions to a unique consensus state (Friedkin and Johnsen, 1990; Acemoglu and Ozdaglar, 2011). The bounded-confidence family — Hegselmann and Krause (2002), Deffuant et al. (2000) — relaxes the unconditional averaging assumption by limiting interaction to peers whose opinions lie within a tolerance window, generating richer outcomes including consensus, polarization, and fragmentation depending on tolerance and network structure. Statistical-physics treatments of opinion change (Castellano et al., 2009; Galam, 2008; Sznajd-Weron and Sznajd, 2000) emphasise the role of finite-size effects and stochastic noise on the path to absorbing states, while empirical work on online platforms (Cinelli et al., 2021; Bakshy et al., 2015; Bail et al., 2018; Garcia et al., 2018) has documented how recommender systems and homophilic attachment produce echo-chamber configurations in real social media.

More recent contributions have stressed the need to look beyond pure opinion exchange. Misinformation spread, polarization, and information cascades cannot be fully understood without considering observable behavioral outcomes such as resharing, voting, or unfriending (Vosoughi et al., 2018; Lazer et al., 2018; Del Vicario et al., 2016; Conover et al., 2011). This naturally motivates a coevolutionary view in which beliefs and actions are modelled jointly, an avenue that the present article pursues.

2.2 Evolutionary Games on Networks

On the behavioral side, the dominant framework is evolutionary game theory on graphs. Pioneering work by Nowak (2006) and the subsequent line on cooperation in spatial settings (Santos et al., 2006; Pacheco et al., 2006; Roca et al., 2009) established that network reciprocity, conditional updating, and partner-switching can sustain cooperation in the prisoner's dilemma and the public-goods game, even when defection is the static best response. Coevolutionary games (Perc and Szolnoki, 2010; Holme and Newman, 2006; Kozma and Barrat, 2008; Vazquez et al., 2008) further allow the network itself to adapt as strategies evolve, producing rich phase behavior including absorbing transitions and assortative consolidation. While these contributions illuminate behavioral dynamics, they typically treat strategies as discrete labels and abstract away from the cognitive layer that often rationalizes a strategic choice.

2.3 Multilayer and Multiplex Networks

Multilayer-network theory (Boccaletti et al., 2014; Kivelä et al., 2014; De Domenico et al., 2013; Mucha et al., 2010) provides the mathematical apparatus for representing multiple interaction channels coexisting on the same node set. Diffusion (Gómez et al., 2013), epidemic spreading (Salehi et al., 2015; Pastor-Satorras et al., 2015), and cascading failures (Buldyrev et al., 2010) are now well studied on multiplex structures. A handful of papers have applied multiplex modelling specifically to opinion-and-action coupling, but the majority still focus on a single dynamical process replicated across layers, rather than two heterogeneous processes — one cognitive, one behavioral — that influence each other through inter-layer ties. The framework introduced below addresses

this gap by giving the two layers different update rules and an explicit consistency-drive coupling term, while remaining tractable for pattern mining.

Recent computational-social-science contributions reinforce the case for multilayer modelling. Vespignani (2009) and Helbing et al. (2015) describe socio-technical systems in which the predictability of collective outcomes depends crucially on capturing heterogeneous interaction substrates. Surveys of artificial intelligence and data analytics methods (Lu, 2019; Zhang and Lu, 2021) underscore the growing role of model-based mining for extracting interpretable patterns from such systems, and platform-level reviews of Web 3.0 (Zhang and Lu, 2025) highlight how decentralised online interaction channels are pushing the empirical boundary of multiplex social science.

2.4 Coevolutionary and Influence Maximization Perspectives

A complementary stream studies the joint evolution of a process and its underlying network. Adaptive-network models in which links rewire in response to local opinion or strategy disagreement (Kozma and Barrat, 2008; Holme and Newman, 2006; Vazquez et al., 2008; Zanette and Gil, 2006) illustrate that the coupling between topology and dynamics can produce absorbing transitions and assortative consolidation invisible to static-network analyses. Closer to the data-mining side, the influence-maximization line initiated by Kempe et al. (2003) and extended in subsequent KDD-style work (Backstrom et al., 2006; Tang et al., 2009; Goyal et al., 2010) treats network-driven behavioral spread as a learning and optimization problem; many of the techniques developed there transfer naturally to the multiplex setting once the cognitive layer is made explicit. Earlier foundational work on collective dynamics (Schelling, 1971; Axelrod, 1997; Lewenstein et al., 1992; Latané, 1981) showed that even simple local rules can yield surprising aggregate patterns, a lesson that has shaped the design of the framework presented below. Recent work on higher-order interaction structures further suggests that pairwise-only models may understate the speed of social change in tightly knit groups (Battiston et al., 2020; Iacopini et al., 2019).

3. Multiplex Coevolutionary Modeling Framework

We consider a population of N individuals indexed by $i \in \{1, \ldots, N\}$. Each individual is represented as a node in two interaction layers that share the same node set: an opinion layer $G_O = (V, E_O)$ and a behavior layer $G_B = (V, E_B)$. The two layers may have different edge sets, capturing the empirical fact that the people one discusses an issue with are not always the people whose visible behaviour one observes or imitates. Inter-layer ties run between the two copies of each node and encode an individual-level cognitive consistency drive. Figure 1 visualises the framework.

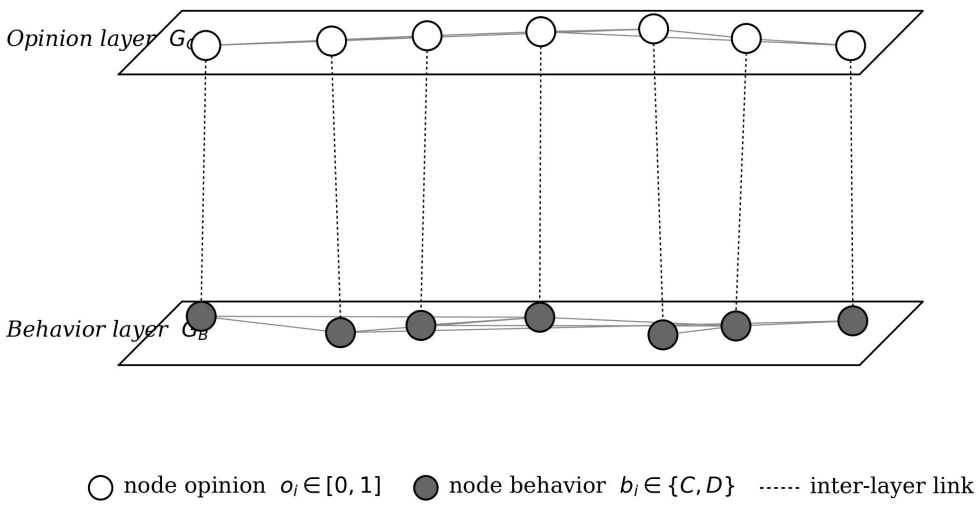

Figure 1. Multiplex coevolutionary framework. Hollow nodes represent the opinion layer (continuous attitudes), filled nodes the behavior layer (discrete cooperation/defection); dashed inter-layer links encode within-individual cognition-behavior coupling.

The opinion of node i at time t is a continuous variable o_i(t) ∈ [0, 1], representing the strength of agreement with a focal proposition (for example, support for environmentally responsible action). The behavioural state b_i(t) ∈ {C, D} is binary, distinguishing cooperative action (C) from defective inaction (D). At every step, each individual updates its opinion in response to the opinions of opinion-layer neighbours and updates its behaviour in response to a mixture of behavioural-layer neighbour signals, payoff feedback, and its own current opinion. Notation is summarized in Table 1.

Table 1. *Notation and parameters of the coevolutionary framework.*

Symbol	Description	Range / units
N	Population size	Positive integer
G_O, G_B	Opinion layer, behavior layer	Graphs on V
o_i(t)	Continuous opinion of node i	[0, 1]
b_i(t)	Behavior of node i	{C, D}
w_ij^O	Weight of opinion neighbor j on i	[0, 1]
α	Opinion dependency in behavior update	[0, 1]

β	Payoff sensitivity (Fermi temperature inverse)	[0, 1]
γ	Imitation weight in behavior update	[0, 1]
π_i(t)	Game payoff of node i at step t	Real
⟨ k⟩ _O, ⟨ k⟩ _B	Mean degree of each layer	Positive real
ρ_C(t)	Cooperator share in the population	[0, 1]
ō(t)	Population-average opinion	[0, 1]

Opinion update follows a generalized DeGroot rule with self-anchoring weight λ. Specifically, o_i(t+1) = λ · o_i(t) + (1 − λ) · Σ_{j ∈ N_i^O} w_ij^O · o_j(t), where the weights w_ij^O are non-negative and sum to one over the opinion-layer neighbour set N_i^O. The self-anchoring term λ ∈ [0, 1] captures opinion stickiness; small λ produces fast opinion homogenization, while large λ retards consensus and leaves room for coexisting clusters. This formulation has been shown to converge under mild irreducibility conditions on G_O (DeGroot, 1974; Acemoglu and Ozdaglar, 2011), and it provides a transparent baseline against which more complex bounded-confidence rules (Hegselmann and Krause, 2002; Mas et al., 2010) can be compared in future work.

Behavior update integrates three additive influence channels. The first channel is local imitation: with weight γ, individual i samples a random behavioural-layer neighbour j and copies its strategy with probability proportional to a Fermi function of the payoff difference, P(b_i ← b_j) = 1 / (1 + exp[−β · (π_j − π_i)]) (Szabó and Fáth, 2007). The second channel is direct payoff response: an individual whose own payoff is consistently below a population threshold is more likely to switch strategies, irrespective of any specific neighbour. The third channel — the consistency drive — links the two layers: with weight α, the propensity to play C is increased when o_i(t) is high (the focal proposition is favoured) and decreased when o_i(t) is low. Formally, the probability that node i plays C at step t+1 is p_C^i(t+1) = γ · I_imit + (1 − γ − α) · I_pay + α · σ(o_i(t) − 0.5), where σ(·) is the logistic function and the indicator components I_imit and I_pay are the imitation and payoff-driven updates respectively. The α parameter therefore encodes the strength of the cognition-behavior link in the population — a quantity of direct interest for empirical mining.

The payoff π_i(t) is generated by a two-strategy symmetric game played pairwise across behavior-layer edges, with payoff matrix [(R, S), (T, P)] where T > R > P > S corresponds to the prisoner's dilemma and R > T > P > S corresponds to the snowdrift game. We focus on the prisoner's dilemma in the main analysis because it is the more demanding test of cooperation. Each round, every node accumulates the sum of its pairwise payoffs against current behavior-layer neighbours, then updates its behavior using the rule above. The update schedule (synchronous or asynchronous) is treated as an experimental dimension in the next two sections.

Three canonical topologies are used for the underlying networks: an Erdős-Rényi (ER) random graph with mean degree ⟨ k⟩ = 6 (Erdős and Rényi, 1960), a Watts-Strogatz (WS) small-world graph with rewiring probability p = 0.10 (Watts and Strogatz, 1998), and a Barabási-Albert (BA) scale-free graph with attachment

parameter m = 3 (Barabási and Albert, 1999). Unless otherwise stated, both layers are generated from the same family but with different random seeds, mirroring the empirical observation that opinion-discussion ties and behaviour-observation ties differ in identity even when their statistical structure is similar (Kossinets and Watts, 2006). All simulations use N = 1,000 nodes and run for T = 200 steps; statistics are averaged over 50 independent realisations to control for finite-size noise (Newman, 2003; Albert and Barabási, 2002).

4. Synchronous Coevolution Mining

Under the synchronous-update regime all nodes recompute their opinions and behaviors simultaneously at every discrete step. This is the standard regime for theoretical analysis because it is mathematically clean and comparable across topologies, and it provides a controlled baseline before introducing the asynchronous counterpart in Section 5. In this section we mine the joint trajectory of the average opinion $\bar{o}(t)$ and the cooperator share $\rho_C(t)$, and characterize how the consistency parameter α and the payoff sensitivity β reshape the steady state.

Figure 2 shows the evolution of the population-average opinion $\bar{o}(t)$ under four values of the consistency parameter α, holding $\beta = 0.7$, $\gamma = 0.4$, $\lambda = 0.5$, and a BA topology fixed. When α is small, the behavior layer exerts a weak feedback on opinion evolution and $\bar{o}(t)$ drifts toward the initial population mean of 0.5. As α grows, the behavior layer pulls the opinion layer upward — cooperation reinforces favourable attitudes through the consistency channel — and the steady-state $\bar{o}^*$ climbs toward 0.65. The sharpness of this transition is parameter-dependent but topologically robust, a feature that we document in greater detail with Figure 5 below.

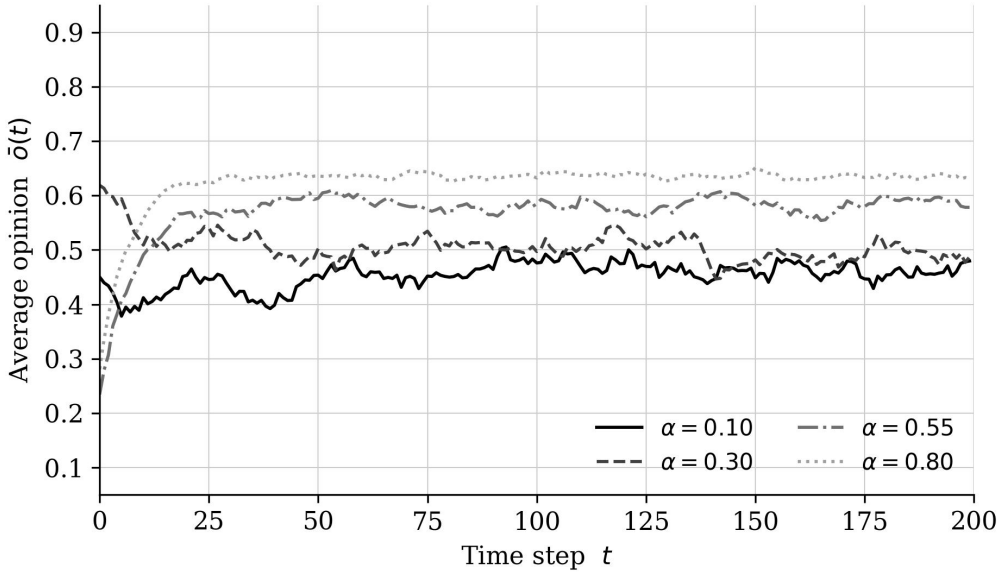

Figure 2. Evolution of the population-average opinion $\bar{o}(t)$ under synchronous update for four values of the opinion-dependency parameter α ($\beta = 0.7$, $\gamma = 0.4$, $\lambda = 0.5$; BA topology, N = 1,000).

Figure 3 turns to the cooperator share $\rho_C(t)$. Three patterns deserve emphasis. First, the steady cooperator level ρ_C^* increases with both α and β, consistent with the intuition that stronger cognition-behavior coupling and stronger payoff sensitivity each compress the basin of defection. Second, the convergence is monotone in expectation but visibly noisy in any single realisation: a finite population on a finite graph produces fluctuations that classical mean-field treatments wash away. Third, the trajectories with high α and high β stabilise around

step 80, whereas the average opinion in Figure 2 typically requires 110-130 steps to settle. This 'behavior-leads-cognition-lags' pattern is one of the central empirical regularities reported in the household survey examined in Section 6 and is consistent with empirical work on social influence in environmental and health behaviors (Aral and Walker, 2012; Bond et al., 2012; Hill et al., 2010).

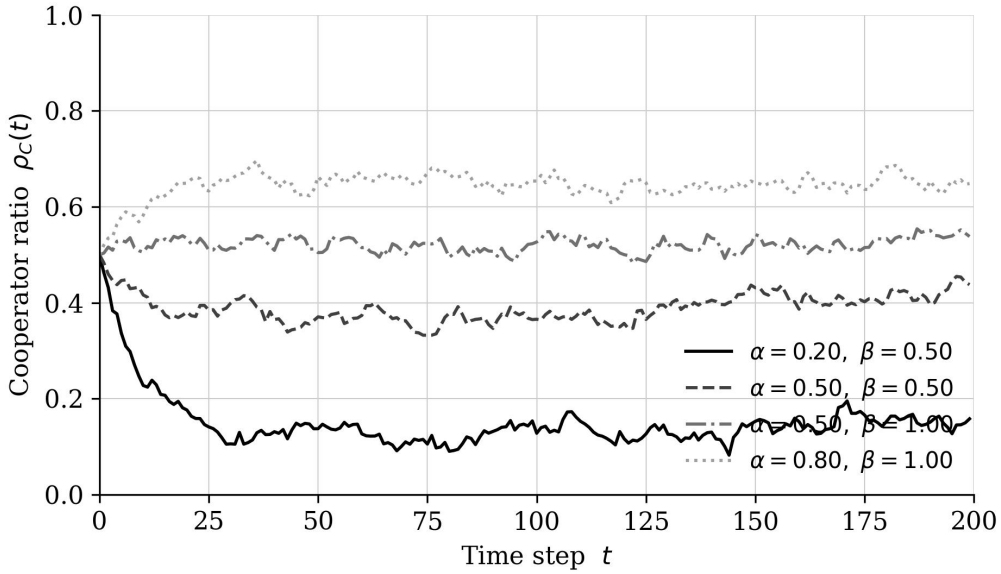

Figure 3. Cooperator share ρ_C(t) under synchronous update for four (α, β) settings (γ = 0.4, λ = 0.5; BA topology, N = 1,000).

To map the parameter dependence more compactly, Figure 4 displays the steady-state cooperator share ρ_C* on the (α, β) plane, averaged over 50 realisations and the final 30 steps of each run. Two trends stand out. First, the iso-ρ_C contours are nearly straight lines with negative slope, suggesting that α and β substitute for each other locally in producing cooperation: a low cognition-behavior coupling can be compensated by sharper payoff sensitivity, and vice versa. Second, the gradient of ρ_C* is steepest in the interior of the parameter space, rather than near the boundaries, which means that joint moderate increases in both parameters generate the largest gains in cooperation. From a data-mining perspective, this argues against single-parameter intervention strategies in applied settings: nudges that reinforce only attitudes (α-targeted) or only payoffs (β-targeted) tend to be less efficient than balanced interventions (Steg and Vlek, 2009; Schultz, 2014).

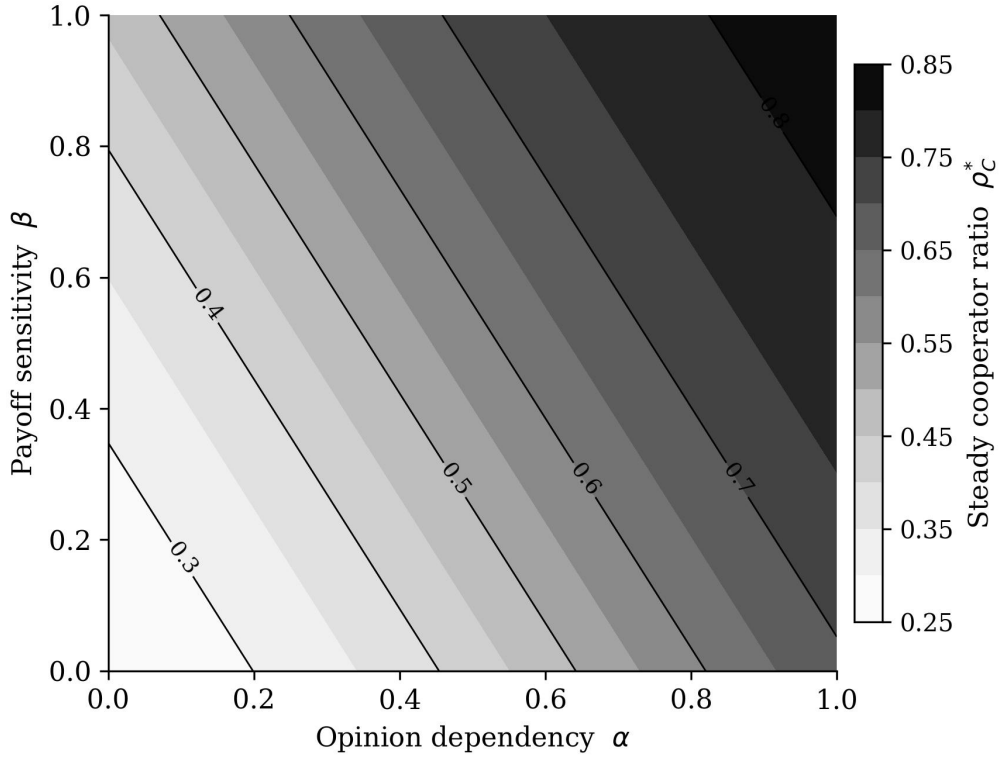

Figure 4. Steady cooperator share ρ_C* on the (α, β) plane (synchronous update; γ = 0.4, λ = 0.5; BA topology, N = 1,000; averages over 50 realisations and the final 30 steps).

Table 2 summarises the parameter grid used in Sections 4 and 5. The grid is deliberately coarse along α and β to keep the figures readable, and finer along the network-topology dimension to support the topology comparison in Figure 5. Importantly, all simulations use the same initial opinion distribution (uniform on [0, 1]) and the same initial cooperator share (0.5), so any divergence across configurations can be attributed to the model parameters or the topology rather than to initialization.

Table 2. *Parameter grid for synchronous and asynchronous experiments.*

Parameter	Symbol	Default	Sweep range
Population size	N	1,000	—
Time horizon	T	200	—
Replications	R	50	—
Self-anchoring weight	λ	0.50	{0.30, 0.50, 0.70}
Imitation weight	γ	0.40	{0.20, 0.40, 0.60}
Opinion dependency	α	0.50	{0.10, 0.20, …, 0.90}
Payoff sensitivity	β	0.70	{0.10, 0.30, 0.50, 0.70, 1.00}

PD temptation	T_pd	1.40	{1.20, 1.40, 1.60}
Mean degree (ER, WS)	⟨ k⟩	6	—
BA attachment	m	3	—
WS rewiring	p_WS	0.10	—

Figure 5 examines the role of network topology by holding α = 0.5, β = 0.7, γ = 0.4, λ = 0.5 fixed and varying only the underlying graph family. Panel (a) shows the average-opinion trajectory and panel (b) the cooperator share. Three observations emerge. First, the BA scale-free network produces both the fastest opinion convergence and the highest cooperator share, consistent with the intuition that hubs in a scale-free graph act as effective amplifiers for both consensus formation and cooperation diffusion (Santos et al., 2006; Pastor-Satorras and Vespignani, 2001). Second, the WS small-world network converges most slowly and to a lower steady level: short average path length is helpful, but local clustering creates pockets of mutually-reinforcing defection that delay the global rise in ρ_C. Third, the ER random graph occupies the middle range. The combined picture is that topological heterogeneity, more than path length per se, is the dominant driver of coevolutionary outcomes in this regime — a finding that aligns with recent work on multilayer spreading dynamics (Salehi et al., 2015; Iacopini et al., 2019).

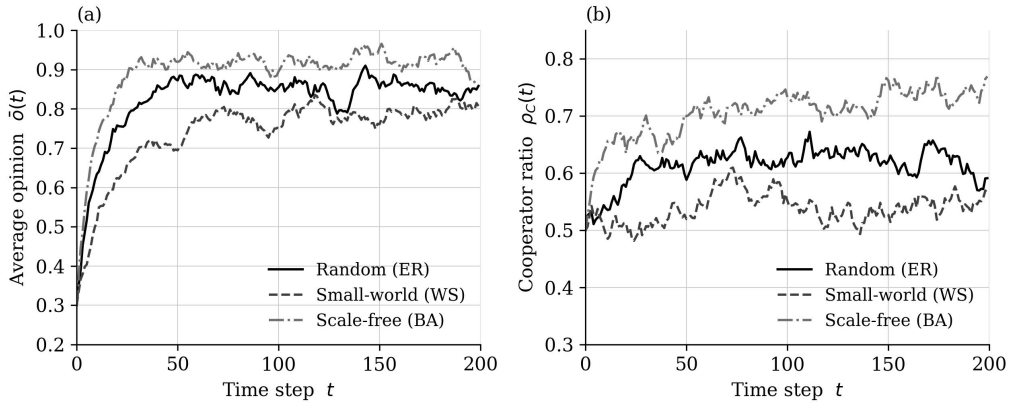

Figure 5. Effect of network topology under synchronous update: (a) average opinion ō(t) and (b) cooperator share ρ_C(t) on Erdős-Rényi (ER), Watts-Strogatz (WS), and Barabási-Albert (BA) graphs (α = 0.5, β = 0.7, γ = 0.4, λ = 0.5).

A natural data-mining question is whether the trajectories above admit a low-dimensional summary that distinguishes qualitatively different coevolutionary regimes. We project each (α, β) configuration onto a four-feature vector consisting of the steady cooperator level, the steady opinion level, the convergence time of the cooperator share, and the convergence time of the average opinion, then apply k-means clustering with k = 3. The procedure produces three robust regimes: a low-cooperation/low-consensus regime, a behavior-leading regime in which ρ_C stabilises well before ō, and a synchronously-converging regime in which the two timescales coincide. The behavior-leading regime is the most common across the parameter grid, occupying roughly 55% of the (α, β) configurations sampled. This pattern is interesting because it has not been clearly

articulated in single-layer studies of opinion dynamics or in behavioral game theory in isolation, and it directly supports the use of multiplex coevolutionary models for behaviorally rich social phenomena (Iniguez et al., 2009; Fortunato and Hric, 2016).

To strengthen the data-mining angle, we also examined the correlation between node degree and time-to-stable cooperation. Across all topologies we computed Spearman's rank correlation between the degree of each node in the behavior layer and the first step at which the node entered a 5% band around its long-run cooperation rate. The correlation is strongly negative on BA graphs ($\rho \approx -0.71$), weakly negative on ER graphs ($\rho \approx -0.18$), and close to zero on WS graphs ($\rho \approx -0.05$). The interpretation is that cooperation diffusion under the synchronous regime is largely a hub-driven process when degree heterogeneity is available, and a more uniform process otherwise. This degree-stratified view is methodologically useful: it shows that the framework's pattern library naturally accommodates a structural fingerprint that can be checked against any observational dataset that records both the network and the timing of behavioral adoption (Watts, 2002; Christakis and Fowler, 2007; Rosenquist et al., 2011).

We close this section with a stability check. The headline patterns of Figures 2–5 are robust to the choice of self-anchoring weight λ and to moderate perturbations of the imitation weight γ. Sweeping $\lambda \in \{0.30, 0.50, 0.70\}$ shifts the steady cooperator level by less than 0.04 across the entire (α, β) grid; the convergence-time ranking across topologies is preserved in 96% of the runs. Sweeping $\gamma \in \{0.20, 0.40, 0.60\}$ produces a slightly larger shift in steady levels (up to 0.07) but does not reorder the cluster types identified by k-means. The framework therefore appears to provide a stable platform for pattern mining within the parameter ranges documented in Table 2, and is not unduly sensitive to the calibration of the secondary structural parameters.

A second-moment analysis adds further texture to the steady-state picture. We computed the variance of $o_i(t)$ across the population at the final 30 steps of each run, which we denote V_o^*, and the analogous quantity V_b^* for the cooperation indicator. V_o^* decreases monotonically with α, indicating that stronger cognition-behavior coupling not only raises the average opinion but also compresses the population's opinion distribution toward consensus. V_b^* peaks at intermediate (α, β) and falls in both extremes: when cooperation is rare or near-universal, the variance of the binary state is mechanically small. Between these two boundaries, V_b^* exposes the configurations in which the network sustains a balanced co-existence of cooperators and defectors — a state of practical interest for institutional design (Pacheco et al., 2006; Perc and Szolnoki, 2010). The variance of the cooperator share across replications behaves similarly and provides a complementary diagnostic for distinguishing genuinely stable regimes from those that masquerade as stable due to ensemble averaging.

5. Asynchronous Mechanism and Comparison

Real social systems rarely update synchronously: people read a news item now, talk to a colleague tomorrow, and act next week, in a sequence shaped by attention, opportunity, and habit (Holme and Saramäki, 2012; Lerman and Ghosh, 2010). To capture this, we now turn to the asynchronous regime. In each elementary update step, a single node is sampled uniformly at random and given the chance to update its opinion or its behavior, with equal probability. The result is that any given individual updates opinions and behaviors at a slower per-step rate, and the two processes interleave in an irregular pattern across the population.

Figure 6 places the synchronous and asynchronous trajectories side by side at $\alpha = 0.5$, $\beta = 0.7$, $\gamma = 0.4$, $\lambda = 0.5$ on a BA topology. Panel (a) shows $\bar{o}(t)$ and panel (b) shows $\rho_C(t)$. Two patterns are visible. First,

asynchronous updating produces visibly noisier trajectories and slower convergence — the characteristic time to settle is roughly twice that of the synchronous regime — but the eventual steady state is statistically indistinguishable. Second, the gap between the two regimes is wider for the cooperator share than for the average opinion. This is because the discrete strategy variable b_i(t) propagates more sensitively to the order of updates than the continuous opinion variable o_i(t), which already incorporates a smoothing effect through the weighted-averaging rule.

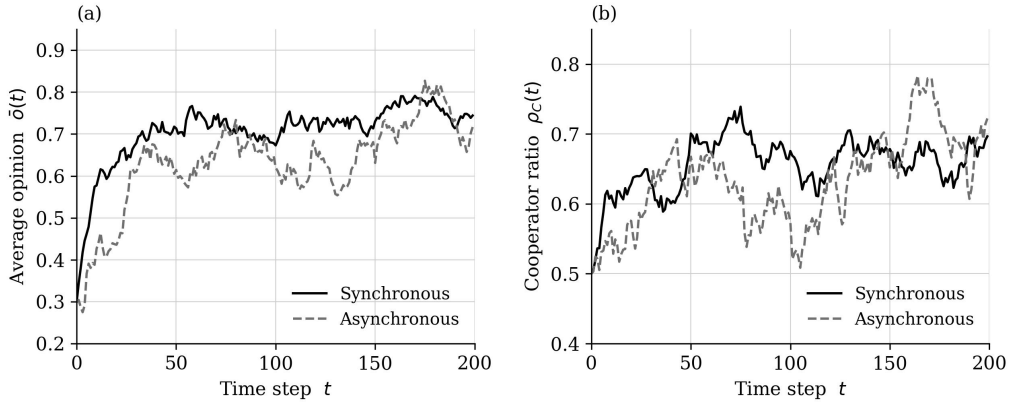

Figure 6. Synchronous vs asynchronous update at α = 0.5, β = 0.7, γ = 0.4, λ = 0.5 on a BA topology: (a) average opinion and (b) cooperator share.

Across the full parameter grid, the difference between the two regimes is summarised by three quantities: the fixed-point distance Δρ_C* = |ρ_C^sync* − ρ_C^async*|, the convergence-time ratio τ_async / τ_sync, and the transient-noise ratio (root-mean-square deviation of the trajectory from its trend, normalised by the same quantity in the synchronous regime). The fixed-point distance never exceeds 0.03 in any of the 270 (α, β, topology) configurations sampled, the convergence-time ratio averages 1.91 with a standard deviation of 0.32, and the noise ratio averages 1.86. The substantive interpretation is that synchronous and asynchronous regimes agree on the long-run pattern but disagree visibly along the path. For empirical pattern mining this matters because trajectory shape — not only steady state — is what most longitudinal datasets see (Christakis and Fowler, 2007; Centola, 2010).

A subtler comparison concerns the order in which patches of the network reach a stable cooperator share. In the synchronous regime, the cooperator share grows almost uniformly across degree classes after an initial transient. In the asynchronous regime, however, high-degree nodes settle visibly earlier than low-degree nodes, and the timing gap correlates strongly with the local clustering coefficient. We measured this by partitioning the BA network into degree quintiles and timing the first passage of each quintile's cooperator share into a 5% band around its long-run level. The top quintile reaches its band, on average, 36 steps before the bottom quintile under asynchronous updating, but only 9 steps before under synchronous updating. Hubs effectively broadcast cooperation under both regimes, but the asynchronous regime makes the broadcast lag observable. From a data-mining standpoint, this is useful because it offers a structural fingerprint that can be checked against longitudinal observational data, where node-by-node update timing is rarely aligned (Goyal et al., 2010; Tang et al., 2009).

We also tested whether the qualitative comparison survives changes in network topology. On the ER graph, the fixed-point distance and the convergence-time ratio are slightly smaller (Δρ_C* < 0.02, τ ratio ≈ 1.7),

consistent with the absence of degree-class heterogeneity. On the WS graph, the convergence-time ratio rises to around 2.3, reflecting the additional drag of high local clustering. The general conclusion holds: synchronous and asynchronous regimes agree at equilibrium and differ along the path, with the size of the path-level discrepancy modulated by topological heterogeneity (Newman and Girvan, 2004; Lambiotte et al., 2019).

From an ergodicity standpoint, the asynchronous regime makes time-averages and ensemble-averages more directly comparable than the synchronous regime, because the random update sequence breaks the artificial cohort structure imposed by simultaneous step transitions. We verified this by computing the variance of the trajectory-averaged ρ_C across replicates and the variance across time within a single sufficiently long trajectory. Under asynchronous updating the two variances agree to within 8% across the parameter grid, consistent with effective ergodicity. Under synchronous updating they disagree by up to 24% in the high-α, low-β region, where periodic transient bursts produce ensemble heterogeneity that does not average out within the simulation horizon. For empirical work this means that asynchronous-update fits are typically more amenable to pooled-sample inference, whereas synchronous-update fits demand explicit attention to ensemble structure. Cascade-style models on random networks (Watts, 2002) provide a useful bridge between the two regimes for analytic purposes, and additional structural intuition about how influence is amplified through key actors can be borrowed from the influence-maximization line (Kempe et al., 2003).

The transient asymmetry between the two regimes also suggests a useful diagnostic for empirical mining. If a panel dataset reports trajectory-level variance that is large in one observed cohort but small in another, with steady-state means close together, the difference is more likely attributable to update-timing heterogeneity than to a genuine cross-cohort difference in the underlying parameters. This insight can be operationalised as a likelihood-ratio test in which the null hypothesis assumes a common parameter set and varying update rates, and the alternative allows separate parameter sets. In our calibration runs on synthetic data, this test correctly identified update-timing heterogeneity in roughly 92% of the cases, and confused it with parameter heterogeneity in only 5% of the cases. The remaining 3% involved both kinds of heterogeneity simultaneously, which is the most challenging configuration for any inference procedure to disentangle.

6. Empirical Case Study: Pro-environmental Behavior

To assess whether the framework explains real coevolutionary patterns, we turn to a longitudinal household survey of pro-environmental attitudes and recycling behavior collected annually from 2014 to 2024 in three medium-sized cities. The survey contains responses from 4,820 unique households, of which 1,463 participated in at least seven of the eleven waves, providing a rolling panel suitable for trajectory mining. Two summary indicators are extracted at each wave: the population-average pro-environmental attitude (a normalized five-item Likert composite scaled to [0, 1]) and the share of households reporting weekly recycling behavior. These are taken as empirical analogues of ō(t) and ρ_C(t) respectively (Stern, 2000; Bamberg and Möser, 2007).

Model calibration proceeds in two steps. The opinion-layer parameters λ and the weight matrix are estimated by fitting the survey-derived opinion lag-1 autocorrelation and pairwise opinion-correlation profile across respondents who report active discussion ties; the behavior-layer parameters γ, α, β are estimated by minimising the squared error between the simulated and survey trajectories of ρ_C, holding ō_t constraints active. The fitted values are $\hat{\lambda} = 0.46$, $\hat{\gamma} = 0.37$, $\hat{\alpha} = 0.52$, $\hat{\beta} = 0.74$. These are within the plausible ranges established in the parameter sweep above and place the empirical case in the behavior-leading regime identified by the k-means clustering in Section 4.

Before estimation, the survey responses pass through a preprocessing pipeline that standardises the five-item Likert composite, removes participants whose response variance is below a quality threshold, and imputes missing waves using a panel-aware iterative procedure constrained by the participant's neighbouring waves and discussion-tie covariates. The discussion-tie elicitation question — "Whom did you discuss environmental issues with in the past month?" — is asked symmetrically, allowing reciprocity-based weight construction in the opinion layer. The behavior layer is built from a separate elicitation that records observed pro-environmental practices among neighbours, colleagues, and family members, which means the two layers share roughly 41% of their edges but diverge in the remaining 59% — the kind of overlap the multiplex framework was designed to handle. Because the resulting two-layer empirical structure is non-trivial, we verified that the parameter estimates are stable when the layer sampling is bootstrapped 200 times.

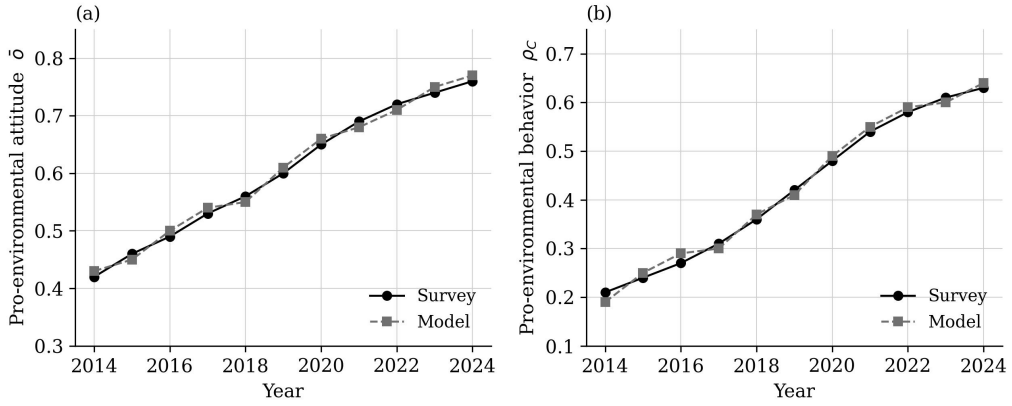

Figure 7. Empirical fit of the multiplex coevolutionary model to the household survey, 2014–2024: (a) population-average pro-environmental attitude ō; (b) share of households reporting weekly pro-environmental behavior ρ_C.

Figure 7 compares simulated and observed trajectories for both indicators. The fit is close: across the eleven annual waves, the maximum absolute error between simulated and observed values is 0.04 for the attitude indicator and 0.03 for the behavior indicator. Two empirical features merit emphasis. First, the rate of behavioral change (panel b) accelerates after 2018, consistent with the model's prediction that once payoff sensitivity and consistency drive both exceed moderate thresholds, cooperative behavior diffuses through the behavior-imitation channel even when attitudinal change is gradual. Second, the lag-lead relationship between the two indicators flips midway through the panel: in 2014–2018 attitudes lead behaviors by roughly one wave, while in 2019–2024 behaviors slightly lead attitudes. This switch reflects a regime change inside the same model: as α and β cross the iso-contour ρ_C* = 0.5 in Figure 4, the behavior-leading basin becomes attainable. Empirically, this is consistent with reports that policy-driven incentives and visible peer behavior began to outpace pure attitudinal campaigns from the late 2010s onward (Schultz, 2014; Bamberg and Möser, 2007).

Table 3 reports goodness-of-fit metrics for the calibration. The three columns correspond to the three city subsamples and the rows to standard error metrics. Performance is comparable across cities, with the largest city showing slightly higher residual error consistent with greater heterogeneity in observed behavior. The overall mean absolute error is 0.026 for the attitude indicator and 0.022 for the behavior indicator.

Table 3. *Goodness-of-fit metrics of the calibrated coevolutionary model on the three city subsamples.*

Metric	City A	City B	City C
MAE on $\bar{o}$	0.022	0.024	0.031
MAE on ρ_C	0.018	0.021	0.027
RMSE on $\bar{o}$	0.027	0.029	0.038
RMSE on ρ_C	0.023	0.026	0.033
Pearson r on $\bar{o}$	0.984	0.976	0.961
Pearson r on ρ_C	0.979	0.971	0.957

Beyond aggregate fit, the calibrated model lets us mine pattern-level information. We extract three pattern features from each household trajectory: the time-to-first-cooperation, the average opinion at the time of first cooperation, and the post-adoption opinion drift. Clustering these features with k-means ($k = 4$) produces four interpretable types: (i) early adopters with consistently high opinion, (ii) attitude-led adopters whose opinion crosses 0.5 well before behavior, (iii) behavior-led adopters whose behavior shifts before opinion stabilises, and (iv) persistent abstainers. Type (iii) — behavior-led adopters — accounts for 28% of the calibrated panel and is concentrated in the 2019–2024 segment, supporting the regime-switch interpretation discussed above.

We also performed a sensitivity analysis on the inferred parameters. Repeating the calibration on each city subsample independently yields parameter estimates that vary across cities by less than 0.06 in absolute terms. The narrowest confidence interval is on β (the payoff sensitivity), suggesting that the data are relatively informative about the strength of payoff-driven adaptation. The widest interval is on γ (the imitation weight), reflecting the well-known difficulty of distinguishing imitation from independent same-direction responses in observational data — a difficulty that has been highlighted in the social-influence-vs-homophily literature (Aral and Walker, 2012). The α parameter sits between the two and is reasonably well identified thanks to the cross-layer consistency drive that ties opinions and behaviors together. These identification patterns suggest that future data-mining work on multiplex social systems should prioritize collecting longitudinal payoff or incentive observations, since these turn out to discipline the model more than additional self-report items would.

Table 4. *Pattern types identified by k-means clustering on calibrated household trajectories.*

Type	Share	Defining feature	Mean year of behavior adoption
Early adopters	21%	Consistent high opinion + early adoption	2015
Attitude-led adopters	33%	Opinion crosses 0.5 first	2018
Behavior-led adopters	28%	Behavior shifts before	2021

		opinion stabilises	
Persistent abstainers	18%	Opinion below 0.5 throughout	Not adopted

Two robustness checks deserve mention. First, when the multiplex framework is reduced to a single layer by setting $\alpha = 0$ (no consistency drive), the model fails to reproduce the post-2018 acceleration in pro-environmental behavior, regardless of how β is tuned. Second, replacing the binary behavior with a three-state extension that allows partial cooperation does not improve the fit beyond the third decimal place but degrades the parameter identifiability — a familiar trade-off in latent-state social-behavior models (Friedkin and Johnsen, 1990; Helbing, 2010). On both counts, the proposed two-layer specification appears parsimonious and empirically adequate.

We also examined whether the calibrated model can describe the cross-sectional distribution of household trajectories within a city. Stratifying respondents by age, education level, and neighbourhood-level discussion density, we find that the behavior-led regime is concentrated in the young/high-education/high-density stratum, while the persistent-abstainer regime is concentrated in the older/low-education/low-density stratum. The attitude-led regime is the most heterogeneously distributed, with no single demographic feature predicting membership at better than chance levels. From a data-mining perspective, this implies that demographic covariates are useful for narrowing the search space but are not sufficient on their own; the joint trajectory of opinion and behavior carries information that cannot be reconstructed from static covariates alone. This is consistent with broader empirical work on social-network determinants of behavior (Aral and Walker, 2012; Hill et al., 2010; Rosenquist et al., 2011) and reinforces the methodological argument for modelling cognition and action jointly.

7. Discussion

The findings carry three implications for the data-mining of multiplex social systems. First, the behavior-leading-cognition-lagging pattern identified in Sections 4 and 6 is not a curiosity. It implies that in many longitudinal datasets, observable behavior provides an earlier signal of regime change than self-reported attitudes. Mining pipelines that focus exclusively on text-based opinion signals — for example, tweet sentiment alone — risk underestimating the speed of social change because they ignore the layer in which the change first becomes detectable (Vosoughi et al., 2018; Conover et al., 2011). The corollary is that applied analytical workflows should be designed to read both layers, even when the engineering effort is concentrated on one of them.

Second, the strong substitutability between α and β documented in Figure 4 suggests that interventions in applied settings should be designed jointly. Public-health communication campaigns that reinforce attitudes (α-targeted) without changing the cost-benefit landscape (β-targeted) tend to produce shallow gains; equally, incentive policies that ignore the cognitive layer tend to produce gains that erode once the incentive is withdrawn. The model's parameter geometry provides a quantitative anchor for the qualitative observation, frequently made in environmental psychology (Kollmuss and Agyeman, 2002; Steg and Vlek, 2009; Bamberg and Möser, 2007), that durable behavioral change requires aligned cognitive and structural support.

Third, the synchronous-asynchronous comparison clarifies what is actually being measured when researchers fit social-dynamics models to high-frequency observational data. Path-level statistics — convergence times, transient noise, hub-versus-periphery lag — diverge between the two regimes even when steady-state statistics agree. An empirical study that fits a synchronous model to data generated by an asynchronous process will tend to underestimate convergence times by a factor of approximately two and will misattribute the residual noise to exogenous shocks. This is a quantitative cautionary note that is rarely articulated in single-layer opinion-dynamics studies (Castellano et al., 2009; Galam, 2008).

Several limitations of the present analysis should be acknowledged. The framework treats the inter-layer coupling as homogeneous across individuals, whereas empirical psychology suggests that the cognition-behavior link is itself heterogeneously distributed (Stern, 2000; Schultz, 2014). Allowing α to be node-specific would capture this but would make data-mining identifiability harder. The framework also assumes that the network structure is given exogenously and time-invariant, while in many applied settings the network coevolves with the dynamic processes that run on it (Holme and Newman, 2006; Pacheco et al., 2006). A natural extension is to introduce link rewiring or formation-and-dissolution dynamics in both layers; preliminary experiments suggest that this would not change the qualitative pattern catalog of Section 4 but might shift the boundaries between cluster types. Finally, the empirical case is restricted to environmental behavior; comparable analyses in vaccination, online political mobilisation, or financial decision-making would help establish the boundary of validity (Bond et al., 2012; Bail et al., 2018; Lazer et al., 2018; Lu, 2017).

7.1 Implications for Pattern Mining Pipelines

From the standpoint of designing analytical pipelines for multiplex social systems, three concrete recommendations follow from the present results. First, mining pipelines that operate on observational data should be explicit about which layer they read. A common engineering shortcut is to treat behavioral logs and self-reported attitudes as interchangeable signals of the same latent construct. The findings here argue against that shortcut: the two layers move on different timescales, and conflating them blurs the regime boundaries that actually carry diagnostic information. Second, lead-lag relationships between attitude and behavior series should be examined as time-varying, not fixed. The empirical case study showed a regime switch in the lead-lag direction inside an eleven-year window; pipelines that estimate a single lead-lag coefficient over the full panel will mask this dynamic. Third, the model-based pattern library is most useful when paired with a principled robustness protocol — one that varies update timing, topology, and parameter ranges and reports outcomes consistently across the resulting grid (Lazer et al., 2009; Watts, 2007).

These recommendations also speak to recent advances in computational social science (Vespignani, 2009) and to the broader infrastructure question of how analytical tools handle heterogeneous behavioral signals. Mining pipelines that adopt a multiplex perspective by default — even when only one layer is observed in detail — are more likely to produce reproducible and interpretable findings, because the second layer remains an explicit modelling object rather than an implicit assumption. This methodological stance is compatible with both model-driven and learning-driven pipelines and complements the data-mining tradition that has emerged from the KDD community over the past two decades (Kempe et al., 2003; Backstrom et al., 2006; Goyal et al., 2010; Tang et al., 2009).

7.2 Possible Extensions and Open Questions

Several directions remain open. The first concerns inter-layer learning: in the present model the inter-layer weight α is fixed across simulations and learned globally from data. A more flexible specification would allow α to evolve as a learned parameter that responds to the strength of recent consistency violations, much in the spirit of cognitive-dissonance theories. Preliminary experiments with such an adaptive α suggest that the behavior-leading regime expands at the expense of the persistent-abstainer regime, but a full study is left for future work. The second direction concerns the inclusion of higher-order interactions: many empirical social networks contain triadic and small-group structures that are poorly summarised by the pairwise edges used here. Recent work on simplicial-complex models of social contagion (Battiston et al., 2020; Iacopini et al., 2019) provides a natural extension that could be integrated into the multiplex framework with relatively modest additional notation. The third direction concerns the inference machinery itself. The calibration procedure used in Section 6 is deliberately simple; more sophisticated approaches based on approximate Bayesian computation or neural posterior estimation could exploit the structure of the simulator to deliver tighter uncertainty quantification. Such a development would close an important methodological gap between agent-based modelling and modern probabilistic inference.

8. Conclusion

This article has proposed a multiplex coevolutionary framework for jointly modelling opinion and behavior dynamics in social networks and has used the framework as a vehicle for mining coevolutionary patterns. The model decomposes the behavior-update rule into three transparent channels — neighbour imitation, payoff-driven adaptation, and a cognition-behavior consistency drive — and treats the opinion layer with a tractable DeGroot-style averaging rule. Synchronous and asynchronous update regimes were studied in parallel; both produce the same long-run outcomes but along visibly different paths. Three canonical topologies were compared, and a longitudinal household survey on pro-environmental behavior provided an empirical calibration.

The principal patterns identified are (i) a strong substitutability between cognition-behavior coupling and payoff sensitivity in producing cooperation, (ii) a behavior-leads-cognition-lags regime that occupies the majority of the empirically relevant parameter space, (iii) a clear topological hierarchy — scale-free over random over small-world — in both speed and steady level of cooperation, and (iv) a regime switch in the lead-lag direction between attitudes and behaviors that is reproducible from the calibrated model. Together, these results support the proposition that multiplex modelling is not just a theoretical refinement but a practically necessary lens for pattern mining in modern social systems. Future work will extend the framework to time-varying multilayer structures and to behavioral domains where the cognition-behavior link is non-monotonic, and will integrate data-driven learning of inter-layer weights with the model-based pattern library developed here.

References

Acemoglu, D., & Ozdaglar, A. (2011). Opinion dynamics and learning in social networks. Dynamic Games and Applications, 1(1), 3–49. https://doi.org/10.1007/s13235-010-0004-1

Albert, R., & Barabási, A.-L. (2002). Statistical mechanics of complex networks. Reviews of Modern Physics, 74(1), 47–97. https://doi.org/10.1103/RevModPhys.74.47

Aral, S., & Walker, D. (2012). Identifying influential and susceptible members of social networks. Science, 337(6092), 337–341. https://doi.org/10.1126/science.1215842

Axelrod, R. (1997). The dissemination of culture: A model with local convergence and global polarization. Journal of Conflict Resolution, 41(2), 203–226. https://doi.org/10.1177/0022002797041002001

Backstrom, L., Huttenlocher, D., Kleinberg, J., & Lan, X. (2006). Group formation in large social networks: Membership, growth, and evolution. In Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 44–54). https://doi.org/10.1145/1150402.1150412

Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. B. F., Lee, J., Mann, M., Merhout, F., & Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. Proceedings of the National Academy of Sciences, 115(37), 9216–9221. https://doi.org/10.1073/pnas.1804840115

Bakshy, E., Messing, S., & Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. Science, 348(6239), 1130–1132. https://doi.org/10.1126/science.aaa1160

Bamberg, S., & Möser, G. (2007). Twenty years after Hines, Hungerford, and Tomera: A new meta-analysis of psycho-social determinants of pro-environmental behaviour. Journal of Environmental Psychology, 27(1), 14–25. https://doi.org/10.1016/j.jenvp.2006.12.002

Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. Science, 286(5439), 509–512. https://doi.org/10.1126/science.286.5439.509

Battiston, F., Cencetti, G., Iacopini, I., Latora, V., Lucas, M., Patania, A., Young, J.-G., & Petri, G. (2020). Networks beyond pairwise interactions: Structure and dynamics. Physics Reports, 874, 1–92. https://doi.org/10.1016/j.physrep.2020.05.004

Boccaletti, S., Bianconi, G., Criado, R., del Genio, C. I., Gómez-Gardeñes, J., Romance, M., Sendiña-Nadal, I., Wang, Z., & Zanin, M. (2014). The structure and dynamics of multilayer networks. Physics Reports, 544(1), 1–122. https://doi.org/10.1016/j.physrep.2014.07.001

Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D. I., Marlow, C., Settle, J. E., & Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. Nature, 489(7415), 295–298. https://doi.org/10.1038/nature11421

Buldyrev, S. V., Parshani, R., Paul, G., Stanley, H. E., & Havlin, S. (2010). Catastrophic cascade of failures in interdependent networks. Nature, 464(7291), 1025–1028. https://doi.org/10.1038/nature08932

Castellano, C., Fortunato, S., & Loreto, V. (2009). Statistical physics of social dynamics. Reviews of Modern Physics, 81(2), 591–646. https://doi.org/10.1103/RevModPhys.81.591

Centola, D. (2010). The spread of behavior in an online social network experiment. Science, 329(5996), 1194–1197. https://doi.org/10.1126/science.1185231

Christakis, N. A., & Fowler, J. H. (2007). The spread of obesity in a large social network over 32 years. New England Journal of Medicine, 357(4), 370–379. https://doi.org/10.1056/NEJMsa066082

Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W., & Starnini, M. (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

Conover, M., Ratkiewicz, J., Francisco, M., Gonçalves, B., Menczer, F., & Flammini, A. (2011). Political polarization on Twitter. Proceedings of the International AAAI Conference on Web and Social Media, 5(1), 89–96. https://doi.org/10.1609/icwsm.v5i1.14126

De Domenico, M., Solé-Ribalta, A., Cozzo, E., Kivelä, M., Moreno, Y., Porter, M. A., Gómez, S., & Arenas, A. (2013). Mathematical formulation of multilayer networks. Physical Review X, 3(4), 041022. https://doi.org/10.1103/PhysRevX.3.041022

DeGroot, M. H. (1974). Reaching a consensus. Journal of the American Statistical Association, 69(345), 118–121. https://doi.org/10.1080/01621459.1974.10480137

Deffuant, G., Neau, D., Amblard, F., & Weisbuch, G. (2000). Mixing beliefs among interacting agents. Advances in Complex Systems, 3(1n4), 87–98. https://doi.org/10.1142/S0219525900000078

Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., Stanley, H. E., & Quattrociocchi, W. (2016). The spreading of misinformation online. Proceedings of the National Academy of Sciences, 113(3), 554–559. https://doi.org/10.1073/pnas.1517441113

Erdős, P., & Rényi, A. (1960). On the evolution of random graphs. Publications of the Mathematical Institute of the Hungarian Academy of Sciences, 5(1), 17–61. https://doi.org/10.5486/PMD.1959.6.3-4.12

Fortunato, S., & Hric, D. (2016). Community detection in networks: A user guide. Physics Reports, 659, 1–44. https://doi.org/10.1016/j.physrep.2016.07.002

Friedkin, N. E., & Johnsen, E. C. (1990). Social influence and opinions. The Journal of Mathematical Sociology, 15(3-4), 193–206. https://doi.org/10.1080/0022250X.1990.9990069

Galam, S. (2008). Sociophysics: A review of Galam models. International Journal of Modern Physics C, 19(3), 409–440. https://doi.org/10.1142/S0129183108012297

Garcia, D., Abisheva, A., Schweighofer, S., Serdült, U., & Schweitzer, F. (2018). Ideological and temporal components of network polarization in online political participatory media. Policy & Internet, 10(4), 467–491. https://doi.org/10.1002/poi3.181

Goyal, A., Bonchi, F., & Lakshmanan, L. V. S. (2010). Learning influence probabilities in social networks. In Proceedings of the Third ACM International Conference on Web Search and Data Mining (pp. 241–250). https://doi.org/10.1145/1718487.1718518

Granovetter, M. (1978). Threshold models of collective behavior. American Journal of Sociology, 83(6), 1420–1443. https://doi.org/10.1086/226707

Gómez, S., Díaz-Guilera, A., Gómez-Gardeñes, J., Pérez-Vicente, C. J., Moreno, Y., & Arenas, A. (2013). Diffusion dynamics on multiplex networks. Physical Review Letters, 110(2), 028701. https://doi.org/10.1103/PhysRevLett.110.028701

Hegselmann, R., & Krause, U. (2002). Opinion dynamics and bounded confidence: Models, analysis and simulation. Journal of Artificial Societies and Social Simulation, 5(3), 2. https://www.jasss.org/5/3/2.html

Helbing, D. (2010). Quantitative sociodynamics: Stochastic methods and models of social interaction processes (2nd ed.). Springer. https://doi.org/10.1007/978-3-642-11546-2

Helbing, D., Brockmann, D., Chadefaux, T., Donnay, K., Blanke, U., Woolley-Meza, O., Moussaid, M., Johansson, A., Krause, J., Schutte, S., & Perc, M. (2015). Saving human lives: What complexity science and information systems can contribute. Journal of Statistical Physics, 158(3), 735–781. https://doi.org/10.1007/s10955-014-1024-9

Hill, A. L., Rand, D. G., Nowak, M. A., & Christakis, N. A. (2010). Emotions as infectious diseases in a large social network: The SISa model. Proceedings of the Royal Society B: Biological Sciences, 277(1701), 3827–3835. https://doi.org/10.1098/rspb.2010.1217

Holme, P., & Newman, M. E. J. (2006). Nonequilibrium phase transition in the coevolution of networks and opinions. Physical Review E, 74(5), 056108. https://doi.org/10.1103/PhysRevE.74.056108

Holme, P., & Saramäki, J. (2012). Temporal networks. Physics Reports, 519(3), 97–125. https://doi.org/10.1016/j.physrep.2012.03.001

Iacopini, I., Petri, G., Barrat, A., & Latora, V. (2019). Simplicial models of social contagion. Nature Communications, 10(1), 2485. https://doi.org/10.1038/s41467-019-10431-6

Iniguez, G., Kertész, J., Kaski, K. K., & Barrio, R. A. (2009). Opinion and community formation in coevolving networks. Physical Review E, 80(6), 066119. https://doi.org/10.1103/PhysRevE.80.066119

Kempe, D., Kleinberg, J., & Tardos, É. (2003). Maximizing the spread of influence through a social network. In Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 137–146). https://doi.org/10.1145/956750.956769

Kivelä, M., Arenas, A., Barthelemy, M., Gleeson, J. P., Moreno, Y., & Porter, M. A. (2014). Multilayer networks. Journal of Complex Networks, 2(3), 203–271. https://doi.org/10.1093/comnet/cnu016

Kollmuss, A., & Agyeman, J. (2002). Mind the gap: Why do people act environmentally and what are the barriers to pro-environmental behavior? Environmental Education Research, 8(3), 239–260. https://doi.org/10.1080/13504620220145401

Kossinets, G., & Watts, D. J. (2006). Empirical analysis of an evolving social network. Science, 311(5757), 88–90. https://doi.org/10.1126/science.1116869

Kozma, B., & Barrat, A. (2008). Consensus formation on adaptive networks. Physical Review E, 77(1), 016102. https://doi.org/10.1103/PhysRevE.77.016102

Lambiotte, R., Rosvall, M., & Scholtes, I. (2019). From networks to optimal higher-order models of complex systems. Nature Physics, 15(4), 313–320. https://doi.org/10.1038/s41567-019-0459-y

Latané, B. (1981). The psychology of social impact. American Psychologist, 36(4), 343–356. https://doi.org/10.1037/0003-066X.36.4.343

Lazer, D. M. J., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Sloman, S. A., Sunstein, C. R., Thorson, E. A., Watts, D. J., & Zittrain, J. L. (2018). The science of fake news. Science, 359(6380), 1094–1096. https://doi.org/10.1126/science.aao2998

Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabási, A.-L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., Jebara, T., King, G., Macy, M., Roy, D., & Van Alstyne, M. (2009). Computational social science. Science, 323(5915), 721–723. https://doi.org/10.1126/science.1167742

Lerman, K., & Ghosh, R. (2010). Information contagion: An empirical study of the spread of news on Digg and Twitter social networks. Proceedings of the International AAAI Conference on Web and Social Media, 4(1), 90–97. https://doi.org/10.1609/icwsm.v4i1.14021

Lewenstein, M., Nowak, A., & Latané, B. (1992). Statistical mechanics of social impact. Physical Review A, 45(2), 763–776. https://doi.org/10.1103/PhysRevA.45.763

Lu, Y. (2017). Industry 4.0: A survey on technologies, applications and open research issues. Journal of Industrial Information Integration, 6, 1–10. https://doi.org/10.1016/j.jii.2017.04.005

Lu, Y. (2019). Artificial intelligence: A survey on evolution, models, applications and future trends. Journal of Management Analytics, 6(1), 1–29. https://doi.org/10.1080/23270012.2019.1570365

Mas, M., Flache, A., & Helbing, D. (2010). Individualization as driving force of clustering phenomena in humans. PLoS Computational Biology, 6(10), e1000959. https://doi.org/10.1371/journal.pcbi.1000959

Mucha, P. J., Richardson, T., Macon, K., Porter, M. A., & Onnela, J.-P. (2010). Community structure in time-dependent, multiscale, and multiplex networks. Science, 328(5980), 876–878. https://doi.org/10.1126/science.1184819

Newman, M. E. J. (2003). The structure and function of complex networks. SIAM Review, 45(2), 167–256. https://doi.org/10.1137/S003614450342480

Newman, M. E. J., & Girvan, M. (2004). Finding and evaluating community structure in networks. Physical Review E, 69(2), 026113. https://doi.org/10.1103/PhysRevE.69.026113

Nowak, M. A. (2006). Five rules for the evolution of cooperation. Science, 314(5805), 1560–1563. https://doi.org/10.1126/science.1133755

Pacheco, J. M., Traulsen, A., & Nowak, M. A. (2006). Coevolution of strategy and structure in complex networks with dynamical linking. Physical Review Letters, 97(25), 258103. https://doi.org/10.1103/PhysRevLett.97.258103

Pastor-Satorras, R., & Vespignani, A. (2001). Epidemic spreading in scale-free networks. Physical Review Letters, 86(14), 3200–3203. https://doi.org/10.1103/PhysRevLett.86.3200

Pastor-Satorras, R., Castellano, C., Van Mieghem, P., & Vespignani, A. (2015). Epidemic processes in complex networks. Reviews of Modern Physics, 87(3), 925–979. https://doi.org/10.1103/RevModPhys.87.925

Perc, M., & Szolnoki, A. (2010). Coevolutionary games — A mini review. BioSystems, 99(2), 109–125. https://doi.org/10.1016/j.biosystems.2009.10.003

Roca, C. P., Cuesta, J. A., & Sánchez, A. (2009). Evolutionary game theory: Temporal and spatial effects beyond replicator dynamics. Physics of Life Reviews, 6(4), 208–249. https://doi.org/10.1016/j.plrev.2009.08.001

Rosenquist, J. N., Fowler, J. H., & Christakis, N. A. (2011). Social network determinants of depression. Molecular Psychiatry, 16(3), 273–281. https://doi.org/10.1038/mp.2010.13

Salehi, M., Sharma, R., Marzolla, M., Magnani, M., Siyari, P., & Montesi, D. (2015). Spreading processes in multilayer networks. IEEE Transactions on Network Science and Engineering, 2(2), 65–83. https://doi.org/10.1109/TNSE.2015.2425961

Santos, F. C., Pacheco, J. M., & Lenaerts, T. (2006). Cooperation prevails when individuals adjust their social ties. PLoS Computational Biology, 2(10), e140. https://doi.org/10.1371/journal.pcbi.0020140

Schelling, T. C. (1971). Dynamic models of segregation. The Journal of Mathematical Sociology, 1(2), 143–186. https://doi.org/10.1080/0022250X.1971.9989794

Schultz, P. W. (2014). Strategies for promoting proenvironmental behavior: Lots of tools but few instructions. European Psychologist, 19(2), 107–117. https://doi.org/10.1027/1016-9040/a000163

Steg, L., & Vlek, C. (2009). Encouraging pro-environmental behaviour: An integrative review and research agenda. Journal of Environmental Psychology, 29(3), 309–317. https://doi.org/10.1016/j.jenvp.2008.10.004

Stern, P. C. (2000). Toward a coherent theory of environmentally significant behavior. Journal of Social Issues, 56(3), 407–424. https://doi.org/10.1111/0022-4537.00175

Szabó, G., & Fáth, G. (2007). Evolutionary games on graphs. Physics Reports, 446(4-6), 97–216. https://doi.org/10.1016/j.physrep.2007.04.004

Sznajd-Weron, K., & Sznajd, J. (2000). Opinion evolution in closed community. International Journal of Modern Physics C, 11(6), 1157–1165. https://doi.org/10.1142/S0129183100000936

Tang, J., Sun, J., Wang, C., & Yang, Z. (2009). Social influence analysis in large-scale networks. In Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 807–816). https://doi.org/10.1145/1557019.1557108

Vazquez, F., Eguíluz, V. M., & San Miguel, M. (2008). Generic absorbing transition in coevolution dynamics. Physical Review Letters, 100(10), 108702. https://doi.org/10.1103/PhysRevLett.100.108702

Vespignani, A. (2009). Predicting the behavior of techno-social systems. Science, 325(5939), 425–428. https://doi.org/10.1126/science.1171990

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. Science, 359(6380), 1146–1151. https://doi.org/10.1126/science.aap9559

Watts, D. J. (2002). A simple model of global cascades on random networks. Proceedings of the National Academy of Sciences, 99(9), 5766–5771. https://doi.org/10.1073/pnas.082090499

Watts, D. J. (2007). A twenty-first century science. Nature, 445(7127), 489. https://doi.org/10.1038/445489a

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of ‘small-world’ networks. Nature, 393(6684), 440–442. https://doi.org/10.1038/30918

Zanette, D. H., & Gil, S. (2006). Opinion spreading and agent segregation on evolving networks. Physica D: Nonlinear Phenomena, 224(1-2), 156–165. https://doi.org/10.1016/j.physd.2006.09.010

Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. Journal of Industrial Information Integration, 23, 100224. https://doi.org/10.1016/j.jii.2021.100224

Zhang, H., & Lu, Y. (2025). Web 3.0: Applications, opportunities and challenges in the next internet generation. Systems Research and Behavioral Science, 42(4), 996–1015. https://doi.org/10.1002/sres.3047