

ARTICLE

Social learning rules and the effectiveness of behavioural policy: an agent-based model

Giuseppe Alessandro Veltri^{1,2}  and Alberto Acerbi²

¹Center for Behavioural and Implementation Science, Yong Loo Lin School of Medicine, National University of Singapore, Singapore and ²Department of Sociology and Social Research, University of Trento, Trento, Italy

Corresponding author: Giuseppe Alessandro Veltri; Email: giuseppe.veltri@unitn.it

(Received 24 November 2025 UTC; revised 8 May 2026 UTC; accepted 19 May 2026 UTC)

Abstract

Behaviour-change interventions unfold in social systems where people learn from others. We develop a stylised agent-based model to examine how four canonical social learning rules – conformist transmission, informational prestige-biased copying, payoff-biased copying and random copying – shape the impact of a simple seeding intervention. Two arms evolve under identical conditions and learning rules, differing only in initial adoption: both start with exactly 5% baseline adopters and the treatment arm additionally seeds 20% of the remaining non-adopters, yielding an exact 25% vs 5% contrast at $t = 0$. Across homogeneous populations, 70/30 mixed ecologies and sweeps over the share of payoff-biased learners, we track adoption trajectories and treatment–control lift; we also vary payoff parameters, prestige informativeness and conformist thresholds in robustness analyses. We find that the same seeding intervention can stall, drift or cascade depending on the learning ecology. In the baseline specification, conformist dynamics exhibit threshold effects that erase treatment gains, prestige-biased and random copying can preserve positive final lift when diffusion remains incomplete and payoff-biased copying mainly changes the diffusion regime rather than preserving large end-point gaps. Robustness checks show that negative payoff premia suppress diffusion, weak or noisy payoff signals can generate treatment advantages, prestige effects depend on how informative prestige is and conformist treatment effects are concentrated in narrow threshold-boundary regions. These results motivate policy heuristics that evaluate interventions relative to local diffusion potential, make successful outcomes visible when payoff cues matter and tailor seeding to the prevailing mix of learning rules.

Keywords: social learning; interventions; diffusion; adoption; cultural evolution

Introduction

Behavioural public policy (BPP) has evolved from an initial focus on individual choice architecture to a broader, system-oriented perspective attentive to spillovers,

heterogeneity and dynamics. Meta-analytic evidence on nudges shows both promise and variability, suggesting that average effects are modest and context-dependent (Mertens *et al.*, 2022). This motivates a shift from asking *whether* an intervention works to *why, for whom and under what social conditions* (Grüne-Yanoff *et al.*, 2024; Hallsworth, 2023; Veltri, 2026; Banerjee and Veltri, 2024) it works, consistent with diffusion theory and networked thresholds (Granovetter, 1978; Watts, 2002; Centola *et al.*, 2018).

Two recent debates make this shift especially important. First, evidence from large-scale nudge units and meta-analyses suggests that behavioural interventions often produce modest and highly context-dependent effects when implemented at scale (DellaVigna and Linos, 2022; Mertens *et al.*, 2022). Second, the i-frame/s-frame debate warns that BPP can become too narrowly focused on individual-level choice architecture, while neglecting the social, institutional and normative systems in which behaviour is embedded (Chater and Loewenstein, 2023; Connolly *et al.*, 2025). Our approach is not to replace individual-level mechanisms with a purely structural account, but to model the meso-level processes through which individual learning and social context interact. Social learning ecologies are precisely such processes: they translate local exposure, social proof, prestige and visible outcomes into population-level diffusion dynamics.

A growing literature in BPP argues that an explicitly *evolutionary* perspective – what Schimmelpfennig and Muthukrishna call the ‘fourth wave: cultural evolutionary behavioural science’ (CEBS) – is a natural next step (Schimmelpfennig and Muthukrishna, 2025). CEBS explains population-level change as the product of transmission dynamics, in which people use structured social learning strategies rather than copy indiscriminately. In particular, conformity (disproportionately copying the majority), prestige bias (copying high-status demonstrators) and payoff bias (copying those who appear more successful), alongside baseline random copying, are expected to dominate in different settings (Kendal *et al.*, 2018; Jiménez and Mesoudi, 2019; Brand *et al.*, 2021). CEBS also foregrounds *within-society heterogeneity*: populations are mosaics of overlapping cultural distributions, so the same policy lever may cascade in one ecology yet stall or backfire in another (Molleman *et al.*, 2014; Berg *et al.*, 2024; Schimmelpfennig and Muthukrishna, 2025).

Importantly for BPP, CEBS also raises governance and ethics questions (e.g. designing with endogenous norm change rather than imposing it) and stresses evaluation over longer horizons and beyond targeted groups – considerations our modelling can surface by tracing spillovers and ceilings over time (Schimmelpfennig and Muthukrishna, 2025). Empirical policy work already shows why these questions matter. Interventions that combine nudges with local social cohesion can outperform standard measures in some neighbourhood contexts (Merkelbach *et al.*, 2021), while information interventions that seed socially recognised ‘gossips’ can diffuse more effectively than random seeding (Banerjee *et al.*, 2019). These studies illustrate that behavioural policy often works through local learning ecologies rather than through isolated individual choice alone. Before presenting our model, the ‘Cultural evolution’ section briefly reviews key concepts from cultural evolution that motivate the specific social learning rules we study.

We develop an agent-based model (ABM) to ask: how do different social learning strategies condition the effectiveness of behaviour-change interventions? We compare four canonical rules – conformist, prestige-biased, payoff-biased and random copying – under both homogeneous and heterogeneous population compositions. Our design introduces an exact initial adoption asymmetry: both arms start from a 5% baseline of adopters, but the treatment additionally seeds 20% of the remaining non-adopters, producing a 20 percentage-point gap. We simulate $N = 400$ agents, use a default horizon $T = 150$ (extended to $T = 1500$ for the slowest mixed ecologies and to $T = 500$ for the random-dominant mixed ecology) and analyse time-resolved adoption trajectories, final uptake and treatment–control lift.

Conceptually, the paper advances three claims. First, the effectiveness of behavioural interventions cannot be understood without reference to the underlying learning ecology. Holding constant the intervention, horizon and mixing structure, the same seeding can fail, meander or cascade purely because people learn differently from one another. Second, payoff-biased copying primarily changes the *diffusion regime* rather than preserving large final treatment gaps: when payoff advantages are strongly positive, both arms can saturate and final lift collapses; when payoff signals are weak or noisy, treatment effects may instead persist. Third, conformist and prestige dynamics are highly conditional: conformity depends on the relation between seeding and thresholds, while prestige becomes directional only when prestige is informative about behavioural payoffs. We use these patterns to derive policy-relevant heuristics about when to prioritise threshold crossing, visible outcomes, credible demonstrators or broad seeding.

The contribution is therefore not to show for the first time that thresholds, central actors or social cues matter. Rather, we place these mechanisms in a common paired treatment–control simulation framework and ask how they change the *shape* of intervention effects over time. This allows us to connect diffusion theory and BPP evaluation: the same endpoint effect can arise from different mechanisms, and the same mechanism can produce early lift, persistent lift or no final lift depending on the ecology. In this sense, the model responds to calls for more mechanistic evidence in behavioural policy (Grüne-Yanoff, 2016) and for BPP research that can reason about how behavioural agents respond to changes in the systems in which they operate (Van Ryzin, 2021; Connolly *et al.*, 2025).

The ‘Cultural evolution’ section reviews the cultural-evolutionary foundations of the learning strategies we model. The ‘Methodology’ section describes the model and simulation design. The ‘Results’ section reports results for homogeneous populations, mixed ecologies, payoff-share sweeps and robustness analyses. The ‘Discussion’ section discusses mechanisms, policy implications and limitations, and the final section concludes the paper.

Cultural evolution

Cultural evolution is a broad, interdisciplinary framework that seeks to explain how socially transmitted behaviours, ideas and norms change over time. Initially building on analogies with biological evolution, it models population-level patterns as the aggregate outcome of individual and social learning (Mesoudi, 2011). The field encompasses

diverse theoretical and methodological approaches: some retain a close analogy with biological evolution – emphasising selection, mutation and inheritance – while others relax the analogy to focus on ecological and cognitive mechanisms that shape how cultural traits emerge, spread and stabilise in specific environments (Sperber, 1996). Despite this variety, cultural evolutionists share the view that cultural change results from differential adoption, retention, transformation and transmission of information, rather than being reducible to individual rational choice or static equilibria. In this sense, cultural evolution offers economists and behavioural scientists a dynamic, process-based account of how preferences, norms and innovations diffuse and stabilise.

A central insight is that people do not learn randomly from others: they deploy *social learning strategies* that bias when, from whom and what they copy (Laland, 2004; Kendal *et al.*, 2018). These strategies (often called *transmission biases*) are typically divided into *content-based* and *context-based* (Boyd and Richerson, 1988). Content-based strategies rely on intrinsic features of traits (e.g. copying a behaviour that appears effective or salient), whereas context-based strategies rely on properties of the social environment independently of a trait's intrinsic features. The latter include, for example, conformity (copying disproportionately the majority), prestige (copying high-status or knowledgeable individuals) and payoff bias (copying those who appear more successful). Such heuristics can generate complex, path-dependent dynamics even in simple environments. At the same time, there is debate about how flexibly and pervasively these strategies are used: empirical studies suggest that social learning is often context-specific, opportunistic and intertwined with communication and reputation management, rather than the outcome of fixed, global rules (Molleman *et al.*, 2014; Morin, 2016; Acerbi, 2019). Clarifying when and why people rely on particular strategies remains essential for linking micro-level behaviour to macro-level diffusion.

ABM provides a methodological bridge between these micro-level assumptions and population-level outcomes (Acerbi *et al.* 2022; Smaldino, 2023). Within cultural evolution, ABMs have illuminated phenomena from the emergence of cooperation and the maintenance of diversity (Henrich and Boyd, 1998; Mesoudi, 2018) to the diffusion of innovations (Mesoudi, 2011). Closest to our aims, ABMs have examined the evolutionary conditions under which social learning strategies arise and how different strategies generate distinct diffusion signatures. Building on this tradition, we operationalise conformity, prestige, payoff bias and random copying as explicit decision rules and test how they condition the effectiveness of seeding interventions in behaviour-change policy.

Methodology

We develop a stylised ABM to examine how the *same* seeding intervention can have different impacts depending on the prevailing *social learning ecology*. The model is deliberately minimal: agents are embedded in a complete-mixing environment (no explicit network structure) and can switch between two behavioural states, non-adoption ($a_i(t) = 0$) and adoption ($a_i(t) = 1$), over discrete time steps. The purpose is not to reproduce a specific empirical case but to identify how different learning mechanisms shape treatment–control differences under common initial conditions.

Model overview

Each simulation compares two arms that evolve under identical conditions and learning rules:

- **Control:** a small baseline share of adopters at $t = 0$.
- **Treatment:** the same baseline adopters *plus* additional seeded adopters at $t = 0$.

All subsequent changes in adoption arise endogenously through social learning.

Agents follow one of four canonical learning strategies drawn from the cultural-evolution literature: conformist transmission, prestige-biased copying, payoff-biased copying and random copying. The baseline model compares these strategies in homogeneous populations, in mixed 70/30 ecologies and in payoff-share sweeps. We then examine robustness to payoff parameters, prestige informativeness and conformist threshold settings.

Simulation conditions

We study four families of learning ecologies.

E1: Homogeneous populations.

All agents use the same strategy (conformity, prestige, payoff-biased copying or random copying). These baseline conditions serve as benchmark signatures that establish each rule's characteristic diffusion profile under identical seeding.

E2: Mixed 70/30 ecologies.

A dominant strategy accounts for 70% of agents, while the remaining three strategies each account for 10%.

E3: Payoff-share sweeps.

We vary the proportion of payoff-biased agents, $p_{\text{payoff}} \in \{0, 0.1, \dots, 1.0\}$, while allocating the remaining share equally across the other three strategies.

E4: Robustness analyses.

We vary (i) the mean and noise of payoff advantages, (ii) the correlation between prestige and payoff advantage and (iii) conformist thresholds, slopes, baseline adoption and seeding rates.

Outcome measures

Let $A(t) = \frac{1}{N} \sum_{i=1}^N a_i(t)$ be the fraction of adopters at time t . We focus on the *time-resolved lift* of treatment relative to control,

$$\Delta A(t) = A^{(\text{trt})}(t) - A^{(\text{ctrl})}(t),$$

as well as summary diagnostics: final lift $\Delta A(T)$, peak lift $\max_t \Delta A(t)$ and its timing, and time-to-threshold metrics in the treatment arm (e.g. the first t such that $A^{(\text{trt})}(t) \geq 0.5$).

Replication design

For each condition, we run paired Monte Carlo replications. Within each replication, agent attributes (strategy assignment, agent-level payoff advantage, prestige and baseline adoption) are held fixed across arms; the only difference is whether the additional seeding is applied at $t = 0$. Unless otherwise stated, we report means across paired replications (for E3 and the robustness sweeps, we use $R = 50$ per grid point).

Simulation set-up

Population, mixing and updating

Each run uses $N = 400$ agents. Unless otherwise stated, the default horizon is $T = 150$ discrete time steps. For the mixed 70/30 ecologies (E2), we extend the horizon to $T = 1500$ when conformity or prestige is the dominant strategy, and to $T = 500$ when random copying is dominant, to make slower convergence visible. The population is *well-mixed*: when an agent samples a demonstrator, it draws from the other $N - 1$ agents. Updates are synchronous: all agents compute $a_i(t + 1)$ from information available at time t .

Initial conditions and intervention

At $t = 0$, we randomly select exactly 5% of agents to be baseline adopters in *both* arms. In the treatment arm, we then seed an additional 20% of agents from the remaining non-adopters. This construction yields an initial adoption level of 5% in control and 25% in treatment, i.e. an exact initial gap of 20 percentage points.

Learning rules

Agents are assigned one learning strategy for the duration of a run.

Conformist transmission.

Let $A(t)$ be the global adoption rate at time t . A conformist agent adopts with probability

$$p_{\text{conf}}(t) = [1 + \exp(-\beta_c(A(t) - m))]^{-1},$$

with baseline parameters $\beta_c = 10$ and $m = 0.5$. The agent then sets $a_i(t + 1) = 1$ with probability $p_{\text{conf}}(t)$ and $a_i(t + 1)$ otherwise. Because $p_{\text{conf}}(t)$ is non-linear in $A(t)$, it implements *disproportionate* copying of the majority: when $A(t) > m$, the majority trait is copied more than proportionally; when $A(t) < m$, it is suppressed.

Prestige-biased copying.

Each agent i is assigned an adoption payoff advantage $\delta_i \sim \mathcal{N}(\bar{\delta}, \sigma_\delta^2)$ with baseline values $\bar{\delta} = 0.25$ and $\sigma_\delta = 0.15$. Prestige scores are then constructed to be positively correlated with δ_i (baseline correlation $\rho = 0.60$) and normalised into strictly positive weights. A prestige-biased agent samples a demonstrator j with probability proportional to prestige and copies the demonstrator's behaviour, $a_i(t + 1) = a_j(t)$. Prestige is therefore an *informational* but imperfect cue to behavioural advantage.

Payoff-biased copying.

Each period, agents receive noisy realised payoffs,

$$\pi_i(t) = a_i(t) \delta_i + \varepsilon_i(t), \quad \varepsilon_i(t) \sim N(0, \sigma_\varepsilon^2),$$

with baseline payoff noise $\sigma_\varepsilon = 0.25$. A payoff-biased agent samples a demonstrator j uniformly and copies the demonstrator's behaviour with probability

$$p_{\text{copy}}(i \leftarrow j) = [1 + \exp(-\beta_p(\pi_j(t) - \pi_i(t)))]^{-1},$$

with baseline sensitivity $\beta_p = 4$. If copying occurs, the learner sets $a_i(t+1) = a_j(t)$; otherwise, it retains $a_i(t)$. Because copying is based on payoff differences rather than on an absorbing adoption state, this rule is bidirectional: agents can adopt or abandon depending on observed payoff differences.

Random copying.

A random copier samples a demonstrator j uniformly and sets $a_i(t+1) = a_j(t)$.

All implementation and figure-generation code is provided in `model32.py` and is designed to reproduce the simulation conditions and figures used in this paper.

Results

Homogeneous populations

We first consider homogeneous populations, where all agents rely on the same social learning rule. This comparison establishes the benchmark diffusion profile associated with each ecology. Figure 1 shows the treatment and control adoption trajectories for the four homogeneous learning ecologies.

Random copying provides the neutral baseline. Because agents simply copy a random peer, there is no intrinsic directional force at the individual level. Yet the initial treatment asymmetry can persist when diffusion remains incomplete. In the baseline simulations, random copying yields the largest persistent final lift among the homogeneous ecologies: control ends at $E[A^{(\text{ctrl})}(T)] = 0.061$, treatment at $E[A^{(\text{trt})}(T)] = 0.267$ and final lift at $E[\Delta A(T)] = 0.205$ (Supplementary Table 2).

Prestige-biased populations now exhibit directional diffusion because prestige is informative about payoff advantage. Treatment still outperforms control, but the final gap is smaller than under random copying: control ends at 0.062, treatment at 0.225 and final lift at 0.163. The treatment effect remains positive, but prestige tends to raise adoption in both arms rather than preserve a large treatment–control separation.

In conformist populations, diffusion again fails to take off under the baseline seeding design. Because both arms remain below the threshold region of the logistic response, adoption collapses towards zero and the final treatment effect is null.

Finally, in payoff-biased populations, the baseline parameterisation drives both arms to full adoption by the end of the horizon. The treatment arm reaches high adoption earlier, but as both arms converge to $A(T) = 1$, the final lift collapses to zero. This should not be interpreted as a mechanical consequence of an absorbing adoption rule: the robustness analyses below show directly that payoff-biased learners do *not* necessarily diffuse the target behaviour when payoff differences are weak, noisy or negative.




Figure 1. Homogeneous population adoption trajectories under four social learning rules. Each panel shows the fraction of adopters over time in the treatment arm (25% adopters at $t = 0$) and the control arm (5% baseline adopters) for conformist, prestige-biased, payoff-biased and random-copying populations. Light lines display individual Monte Carlo runs; dashed lines plot the replication mean trajectory and solid lines the replication median for each arm ($N = 400$, $T = 150$, paired seeds, 100 replications per strategy).

Mixed populations

We next turn to mixed ecologies, where one learning rule dominates (70% of agents) and the remaining 30% is split evenly among the other three. Figure 2 presents the treatment and control adoption trajectories for the four 70/30 mixed ecologies.

The most consequential pattern concerns random- and prestige-dominant mixtures. A 10% minority of payoff-biased learners does not mechanically force all mixed ecologies to saturation. Instead, final lift depends on whether the dominant ecology itself preserves incomplete diffusion.

Conformity-dominant mixtures still stall. Even over the extended horizon ($T = 1500$), both arms remain close to zero and final lift vanishes. This suggests that under the baseline threshold and seeding combination, a small minority of non-conformists is insufficient to push the system across the conformist boundary.

Prestige-dominant mixtures behave very differently. Because prestige is now informational, both arms climb to high adoption over the long horizon: control ends at 0.896 and treatment at 0.931. The final lift is positive but modest (0.034), because informational prestige raises diffusion in both arms.

Payoff-dominant mixtures also converge to saturation in both arms, again erasing end-of-horizon lift. In contrast, random-dominant mixtures preserve the largest residual gap, although the gap narrows when the horizon is extended. With $T = 500$,




Figure 2. Mixed population adoption trajectories for 70/30 compositions of social learning rules. Each panel shows the adoption paths in the treatment arm (25% adopters at $t = 0$) and the control arm (5% baseline adopters) when one rule is dominant (70% of agents) and the remaining 30% is split evenly among the other three strategies. Light lines display individual Monte Carlo runs; dashed lines plot the replication mean trajectory and solid lines the replication median for each arm. Simulations use $N = 400$ agents; the horizon is $T = 1500$ for the conformity- and prestige-dominant panels, $T = 500$ for the random-dominant panel and $T = 150$ for the payoff-dominant panel.

control ends at 0.689, treatment at 0.780 and final lift is 0.091. In this ecology, the treatment advantage persists because background diffusion remains incomplete in control, but longer observation shows gradual catch-up.

We therefore use $T = 1500$ for the conformity- and prestige-dominant panels, $T = 500$ for the random-dominant panel and $T = 150$ for the payoff-dominant panel.

The time-resolved treatment effect, $\Delta A(t)$, again exhibits distinctive temporal signatures across ecologies (Supplementary Material). In homogeneous populations, payoff-biased learning produces a pronounced early pulse that vanishes as both arms saturate, whereas random copying preserves a more stable lift profile and prestige accumulates lift more gradually. In mixed ecologies, the same logic applies: payoff-dominant systems show large early gains but almost no final lift, random-dominant systems preserve the largest residual gap but show gradual catch-up when followed to $T = 500$, prestige-dominant systems build lift slowly and retain a modest end-point advantage and conformity-dominant systems remain close to zero throughout.

Varying the share of payoff-biased learners

We next vary the share of payoff-biased learners in otherwise mixed populations. For each grid value $p_{\text{payoff}} \in \{0.0, 0.1, \dots, 1.0\}$, the remaining $1 - p_{\text{payoff}}$ fraction of agents is split equally among conformist, prestige-biased and random-copying learners. **Figure 3**




Figure 3. Payoff-share sweep. Mean final adoption in control and treatment, $A^{(\text{ctrl})}(T)$ and $A^{(\text{trt})}(T)$, and mean final lift $\Delta A(T)$, as a function of the payoff-biased share p_{payoff} . For each p_{payoff} , the remaining $1 - p_{\text{payoff}}$ share is split equally among conformist, prestige-biased and random-copying learners. Lines/markers show replication means (paired across arms).

reports mean final adoption in treatment and control, together with mean final lift $\Delta A(T)$.

The aggregate sweep reveals a narrow transition region. When $p_{\text{payoff}} = 0$ or 0.1 , both arms remain near baseline. At $p_{\text{payoff}} = 0.2$, adoption rises but remains incomplete, and final lift reaches its maximum in the aggregate sweep ($E[\Delta A(T)] \approx 0.019$). By $p_{\text{payoff}} = 0.3$, both arms are close to saturation (0.990 vs 0.995), and final lift has already compressed again (0.006). Beyond this point, both arms converge towards full adoption and the end-of-horizon gap vanishes.

The main implication is that payoff-biased learners in this model primarily *shift the diffusion regime* rather than preserve large final treatment effects. They can move a system from low adoption to near-complete adoption, but once both arms cross that transition region, the intervention's final advantage collapses. Supplementary Figure 6 shows that this aggregate pattern masks substantial ecological heterogeneity: two-strategy sweeps against conformity, prestige and random copying have very different transition locations and gap profiles.

Robustness to payoff, prestige and conformity parameters

The baseline results above rely on one parameterisation of payoff-biased copying, informational prestige and conformist thresholds. To assess how much of the narrative depends on those specific values, we varied the key parameters governing each mechanism. For readers less familiar with simulation terminology, this section treats the model as a map from assumptions to outcomes: we change one or two parameters at a time and examine whether adoption remains low, reaches saturation in both arms or produces a treatment–control gap.

First, the payoff robustness sweep (Supplementary Figure 7) shows that the effect of payoff-biased learners depends strongly on the payoff environment. When the average adoption payoff premium is negative ($\bar{\delta} = -0.25$), both arms remain at zero regardless of payoff noise. When the premium is strongly positive ($\bar{\delta} \geq 0.25$), both arms saturate and final lift collapses. The largest final lift appears when average payoff differences are weak or neutral but noisy: for $\bar{\delta} = 0$ and payoff-noise SD 0.5, control ends around 0.319, treatment around 0.734 and final lift reaches 0.415.

Second, the prestige robustness sweep (Supplementary Figure 8) shows that prestige effects depend on how informative prestige is about behavioural advantage. Final lift remains positive across the tested prestige–payoff correlations, but its magnitude is not monotonic: in the present parameterisation, lift varies from about 0.117 to 0.233. The key point is not monotonicity but that prestige becomes directional only when prestige carries informational content.

Third, the conformity robustness sweeps (Supplementary Figures 9 and 10) show that the baseline conformist stall is only one point on a broader threshold surface. By a threshold surface, we mean the set of outcomes generated by different combinations of baseline adoption, seeding, threshold location and response steepness. In the threshold–seed grid, most combinations yield zero final lift, but large effects appear near threshold boundaries; for example, with baseline adoption 0.10, seeding 0.30 and threshold 0.40, final lift is 0.975. In the slope–threshold grid, almost all combinations again yield zero lift except a narrow region with a low threshold and very steep conformity response (e.g. $\beta_c = 20$, threshold 0.30, final lift 0.818). In plain terms, the system behaves as if it has a boundary between two regimes: on one side, both arms stall, and on the other side, both arms diffuse. Large treatment effects appear only near that boundary, where the treatment arm can cross the threshold while the control arm does not. The conformist result is therefore best interpreted as a phase-boundary effect rather than a generic property of conformist transmission.

Discussion

Our simulations support three general conclusions. First, the effectiveness of a behavioural intervention is inseparable from the social learning ecology in which it is deployed: the same seeding intervention can fail, drift or cascade depending on whether diffusion is driven mainly by conformity, prestige, payoff cues or random copying. Second, payoff-biased learning should not be understood as a generic source of persistent treatment lift. Instead, it changes the *diffusion regime*: negative payoff signals suppress adoption, weak or noisy payoff differences can create treatment advantages and strongly positive payoffs drive both arms to the same high-adoption state. Third, conformity and prestige are conditional mechanisms: conformity depends acutely on threshold placement relative to seeding, whereas prestige matters only when it carries information about behavioural value.

A first contribution is therefore methodological. By analysing treatment and control trajectories over time, the model shows that social learning ecologies leave distinctive *dynamic fingerprints* in $\Delta A(t)$. Payoff-biased systems exhibit early spikes that compress as both arms converge; prestige-biased systems can build lift gradually when prestige is informative; random copying preserves lift when diffusion remains

incomplete; and conformist systems can be almost inert except near threshold boundaries. These signatures should be interpreted as stylised rather than diagnostic in any strong empirical sense, but they clarify why endpoint-only evaluations can miss where interventions are doing most of their work.

This implication is directly relevant to current debates about the efficacy of behavioural interventions. If effects are modest on average, one response is to ask whether interventions are too weak; another is to ask whether they are being evaluated without sufficient knowledge of the social mechanism through which they spread. Our results support the second interpretation. To infer a local learning ecology, policymakers would need longitudinal uptake data, process evidence about who observes whom and contextual knowledge about which actors are credible, which outcomes are visible and whether the target behaviour is locally rewarded or normative. In this sense, the model does not provide a direct diagnostic test, but clarifies what kinds of evidence would be needed to move from endpoint evaluation to mechanism-sensitive policy design (Grüne-Yanoff, 2016; Van Ryzin, 2021).

A second insight concerns heterogeneity. The mixed ecologies and the payoff-share sweep show that minority payoff-biased learners can matter disproportionately, but not in a simple monotonic way. Final lift is concentrated in transition regions: once enough payoff-biased learners are present to push both arms towards saturation, the end-of-horizon treatment advantage quickly disappears. This is why the aggregate payoff-share sweep peaks only narrowly around $p_{\text{payoff}} \approx 0.2$ and collapses by 0.3, even though total adoption continues to rise.

These findings also clarify how prestige should be interpreted. Informational prestige is a directional but noisy cue to behavioural value. That makes prestige-dominant ecologies substantively interesting for policy: targeting high-prestige demonstrators can raise uptake in both arms, but final treatment lift depends on whether prestige is informative enough to create faster diffusion in treatment without simply saturating the system.

Conformity, meanwhile, should be handled as a threshold-sensitive mechanism. The baseline result illustrates a threshold stall under the chosen seeding design, but the robustness analyses show that the conformist treatment effect is highly parameter-sensitive. This suggests that conformist-heavy settings are not best characterised by a single expected outcome; rather, they are settings in which the intervention designer should worry about where the population sits relative to local tipping points.

These dynamics support conditional policy guidance. Table 1 summarises the implications.

Beyond the specific guidance summarised in Table 1, our results clarify a more general point: in designing behavioural policy interventions, the relevant unit of design is not only the individual decision-maker but the *population-level learning ecology*. Different ecologies imply different design logics. In threshold-sensitive settings, interventions must push populations across a boundary; in prestige-rich settings, they must leverage credible demonstrators; in payoff-oriented settings, they must make outcome signals visible; and in random or weakly structured ecologies, they may need broad rather than highly targeted seeding.

This also places the model between the i-frame and s-frame positions in recent debates about BPP. A social learning ecology is not merely an individual bias, nor is

Table 1. Strategy-specific guidance for intervention design under Model 32

Population type	Recommended intervention	Expected effect
Conformist-heavy	Diagnose local thresholds; alter perceived majorities or seed above the threshold	Threshold-sensitive; often nil unless near a tipping point
Prestige-dominant	Target genuinely informative/prestigious demonstrators	Moderate but variable; raises adoption in both arms
Payoff-oriented	Make outcomes visible and attributable	Regime change and faster diffusion; final lift often collapses at saturation
Mixed/Random	Broad seeding and reinforcement	Persistent lift when background diffusion remains incomplete

it a fully institutional policy instrument. It is a meso-level mechanism linking individual updating to system-level diffusion. The framework therefore suggests one way in which behavioural science can contribute to s-frame policy questions: not by treating nudges as substitutes for structural change, but by analysing how social proof, prestige and payoff visibility mediate the implementation and uptake of broader interventions (Chater and Loewenstein, 2023; Hallsworth, 2023; Connolly *et al.*, 2025).

Finally, our analysis raises governance and ethical considerations. Social learning rules are not neutral policy levers: influencing conformity, prestige or payoff cues shapes which behaviours become normative, whose outcomes are highlighted and whose status is amplified. If prestige is tied to existing inequalities, prestige-based interventions can reinforce them. If payoff cues are selectively visible, payoff-oriented interventions can privilege some experiences over others. Different ethical perspectives would emphasise different constraints. A welfare-oriented view would ask whether manipulating social cues improves outcomes without excessive distributional costs. An agency- or boosting-oriented view would emphasise transparency and citizens' ability to understand and contest the cues being amplified (Hertwig and Grüne-Yanoff, 2017). A democratic or institutional perspective would ask who decides which outcomes count as 'payoffs', whose prestige is amplified and whether conformity-oriented interventions suppress valuable minority behaviours (Lepenies and Malecka, 2018; Fabian and Pykett, 2022). Embedding cultural-evolutionary behavioural science in public policy therefore requires pairing dynamical analysis with normative criteria such as fairness, transparency and participation.

The model remains deliberately simplified and should be read as a proof-of-concept rather than a calibrated account of any specific policy domain. We retain complete mixing, no explicit institutional structure and no endogenous evolution of learning strategies. Even so, the robustness analyses make clear that the main qualitative lesson survives these simplifications: intervention effects depend jointly on how people learn from others, how informative those social cues are and where the population sits relative to the transition boundary between stagnation and diffusion.

Conclusions

The central implication of our results is that the effectiveness of behavioural interventions cannot be evaluated independently of the learning ecology in which they are deployed. The same seeding intervention can produce near-zero final effects under conformity, moderate persistent lift under prestige or random copying, or large total adoption with almost no residual final lift under payoff-biased copying. Diffusion potential is therefore not an intrinsic property of the intervention alone; it is an emergent property of the interaction between seeding and social learning.

A second implication concerns the role of payoff cues. Payoff-biased copying does not simply guarantee adoption. Instead, it changes the diffusion regime conditional on the payoff environment. Negative payoff premia suppress diffusion; weak or noisy payoff cues can generate sizable treatment advantages; strongly positive payoff premia move both arms to the same high-adoption state and therefore erase end-point lift. For policy, this means that making outcomes visible is potentially powerful, but only if those outcomes are genuinely favourable and interpretable by observers.

A third implication concerns robustness. The present results show that several intuitive claims about learning ecologies need to be stated conditionally. Prestige matters only when it is informative; conformity effects are highly sensitive to threshold placement; and final treatment lift can be largest not where diffusion is strongest, but where background diffusion remains incomplete. This reaffirms the value of time-resolved evaluation: interventions may matter most by changing the speed and regime of diffusion, rather than by creating durable end-point gaps.

Future work can extend this framework along several dimensions. First, relaxing the complete-mixing assumption would allow us to explore heterogeneous network structures (e.g. clustered, scale-free or multiplex graphs) and how they interact with different social learning rules. Second, rather than fixing a one-shot 25% vs 5% seeding contrast at $t = 0$, future work could treat seeding magnitude, targeting and timing as design variables. Third, adaptive interventions – in which policymakers update seeding or messaging in response to intermediate outcomes – would bring the simulations closer to real-world iterative programming. Fourth, heterogeneous ecologies without payoff-biased learners, for example mixtures of conformity, informational prestige and random copying only, would help isolate what payoff cues add relative to other forms of social transmission. Finally, linking adoption outcomes to simple cost metrics would make it possible to compare the cost-effectiveness of threshold-focused, prestige-based, payoff-oriented and broad-seeding strategies on a common footing. More broadly, the model points to a BPP agenda that treats social learning as part of the policy environment itself: behavioural interventions may be individual-facing, but their effectiveness often depends on s-frame conditions such as social networks, institutional credibility and the visibility of outcomes.

Supplementary material. To view supplementary material for this article, please visit <https://doi.org/10.1017/bpp.2026.10046>.

Data and code availability. All data and code used in this study are available at the OSF link: <https://doi.org/10.17605/OSF.IO/BPD87>.

References

- Acerbi, A. (2019), *Cultural Evolution in the Digital Age*, Oxford: Oxford University Press.
- Acerbi, A., A. Mesoudi and M. Smolla (2022), *Individual-Based Models of Cultural Evolution: A Step-by-Step Guide Using R*, London: Routledge.
- Banerjee, A., A. G. Chandrasekhar, E. Duflo and M. O. Jackson (2019), 'Using gossips to spread information: theory and evidence from two randomized controlled trials', *The Review of Economic Studies*, **86**(6): 2453–2490. doi:[10.1093/restud/rdz008](https://doi.org/10.1093/restud/rdz008).
- Banerjee, S. and G. A. Veltri (2024), 'Harnessing pluralism in behavioral public policy requires insights from computational social science', *Frontiers in Behavioral Economics*, **3**: 1503793.
- Berg, P. V. D., T. Van Vu and L. Molleman (2024), 'Unpredictable benefits of social information can lead to the evolution of individual differences in social learning', *Nature Communications*, **15**. doi:[10.1038/s41467-024-49530-4](https://doi.org/10.1038/s41467-024-49530-4).
- Boyd, R. and P. J. Richerson (1988), *Culture and the Evolutionary Process*, Chicago: University of Chicago Press.
- Brand, C. O., A. Mesoudi and T. J. H. Morgan (2021), 'Trusting the experts: the domain-specificity of prestige-biased social learning', *PloS One*, **16**(8): e0255346. doi:[10.1371/journal.pone.0255346](https://doi.org/10.1371/journal.pone.0255346).
- Centola, D., J. Becker, D. Brackbill and A. Baronchelli (2018), 'Experimental evidence for tipping points in social convention', *Science*, **360**(6393): 1116–1119. doi:[10.1126/science.aas8827](https://doi.org/10.1126/science.aas8827).
- Chater, N. and G. Loewenstein (2023), 'The i-frame and the s-frame: how focusing on individual-level solutions has led behavioral public policy astray', *Behavioral and Brain Sciences*, **46**: e147. doi:[10.1017/S0140525X22002023](https://doi.org/10.1017/S0140525X22002023).
- Connolly, D. J., G. Loewenstein and N. Chater (2025), 'An s-frame Agenda for behavioral public policy research', *Behavioural Public Policy*, **9**(3): 1–21. doi:[10.1017/bpp.2024.58](https://doi.org/10.1017/bpp.2024.58).
- DellaVigna, S. and E. Linos (2022), 'RCTs to scale: comprehensive evidence from two nudge units', *Econometrica*, **90**(1): 81–116. doi:[10.3982/ECTA18709](https://doi.org/10.3982/ECTA18709).
- Fabian, M. and J. Pykett (2022), 'Be happy: navigating normative issues in behavioral and well-being public policy', *Perspectives on Psychological Science*, **17**(1): 169–182. doi:[10.1177/1745691620984395](https://doi.org/10.1177/1745691620984395).
- Granovetter, M. (1978), 'Threshold models of collective behavior', *American Journal of Sociology*, **83**(6): 1420–1443. doi:[10.1086/226707](https://doi.org/10.1086/226707).
- Grüne-Yanoff, T. (2016), 'Why behavioural policy needs mechanistic evidence', *Economics and Philosophy*, **32**(3): 463–483. doi:[10.1017/S0266267115000425](https://doi.org/10.1017/S0266267115000425).
- Grüne-Yanoff, T., C. Marchionni and T. Nuotio (2024), 'The relevance of mechanisms and mechanistic knowledge for behavioural interventions: the case of household energy consumption', *Economics & Philosophy* **40**, 606–625. doi:[10.1017/S026626712300038X](https://doi.org/10.1017/S026626712300038X).
- Hallsworth, M. (2023), 'A manifesto for applying behavioural science', *Nature Human Behaviour*, **7**(3): 310–322. doi:[10.1038/s41562-022-01501-9](https://doi.org/10.1038/s41562-022-01501-9).
- Henrich, J. and R. Boyd (1998), 'The evolution of conformist transmission and the emergence of between-group differences', *Evolution and Human Behavior*, **19**(4): 215–241.
- Hertwig, R. and T. Grüne-Yanoff (2017), 'Nudging and boosting: steering or empowering good decisions', *Perspectives on Psychological Science*, **12**(6): 973–986. doi:[10.1177/1745691617702496](https://doi.org/10.1177/1745691617702496).
- Jiménez, Á. V. and A. Mesoudi (2019), 'Prestige-biased social learning: current evidence and outstanding questions', *Palgrave Communications*, **5**(1): 1–12. doi:[10.1057/s41599-019-0228-7](https://doi.org/10.1057/s41599-019-0228-7).
- Kendal, R. L., N. J. Boogert, L. Rendell, K. N. Laland, M. Webster and P. L. Jones (2018), 'Social learning strategies: bridge-building between fields', *Trends in Cognitive Sciences*, **22**(7): 651–665. doi:[10.1016/j.tics.2018.04.003](https://doi.org/10.1016/j.tics.2018.04.003).
- Laland, K. N. (2004), 'Social learning strategies', *Learning & Behavior*, **32**(1): 4–14.
- Lepenieş, R. and M. Małecka (2018), 'The Ethics of Behavioural Public Policy', in A. Lever and A. Poama (eds), *The Routledge Handbook of Ethics and Public Policy*, pp. 513–525, Abingdon: Routledge.
- Merkelbach, I., M. Dewies and S. Denktas (2021), 'Committing to keep clean: nudging complements standard policy measures to reduce illegal urban garbage disposal in a neighborhood with high levels of social cohesion', *Frontiers in Psychology*, **12**: 660410. doi:[10.3389/fpsyg.2021.660410](https://doi.org/10.3389/fpsyg.2021.660410).
- Mertens, S., M. Herberz, U. J. J. Hahnel and T. Brosch (2022), 'The effectiveness of nudging: a meta-analysis of choice architecture interventions across behavioral domains', *Proceedings of the National Academy of Sciences*, **119**(1): e2107346118. doi:[10.1073/pnas.2107346118](https://doi.org/10.1073/pnas.2107346118).

- Mesoudi, A. (2011), *Cultural Evolution: How Darwinian Theory Can Explain Human Culture and Synthesize the Social Sciences*, Chicago: University of Chicago Press.
- Mesoudi, A. (2018), 'Migration, acculturation, and the maintenance of between-group cultural variation', *PloS One*, **13**(10): e0205573.
- Molleman, L., P. van den Berg and F. J. Weissing (2014), 'Consistent individual differences in human social learning strategies', *Nature Communications*, **5**. doi:[10.1038/ncomms4570](https://doi.org/10.1038/ncomms4570).
- Morin, O. (2016), *How Traditions Live and Die*, Oxford: Oxford University Press.
- Schimmelpfennig, R. and M. Muthukrishna (2025), 'Cultural evolutionary behavioural science in public policy', *Behavioural Public Policy*, **9**: 652–682. doi:[10.1017/bpp.2022.40](https://doi.org/10.1017/bpp.2022.40).
- Smaldino, P. (2023), *Modeling Social Behavior: Mathematical and Agent-Based Models of Social Dynamics and Cultural Evolution*, Princeton, NJ: Princeton University Press.
- Sperber, D. (1996), *Explaining Culture: A Naturalistic Approach*, Oxford: Blackwell.
- Van Ryzin, G. G. (2021), 'Nudging and muddling through', *Perspectives on Public Management and Governance*, **4**(4): 339–345. doi:[10.1093/ppmgov/gvab021](https://doi.org/10.1093/ppmgov/gvab021).
- Veltri, G. A. (2026), 'Harnessing heterogeneity in behavioural research using computational social science', *Behavioural Public Policy* **10**(2), 382–399. doi:[10.1017/bpp.2023.35](https://doi.org/10.1017/bpp.2023.35).
- Watts, D. J. (2002), 'A simple model of global cascades on random networks', *Proceedings of the National Academy of Sciences*, **99**(9): 5766–5771. doi:[10.1073/pnas.082090499](https://doi.org/10.1073/pnas.082090499).