WORKING PAPERS

N° 1679

October 2025

“Contrarian Motives in Social Learning”

Vasilii Ivanik and Georgy Lukyanov

Toulouse School of Economics

Contrarian Motives in Social Learning*

Vasilii Ivanik
ICEF, NRU HSE
vsivanik@edu.hse.ru

Georgy Lukyanov
Toulouse School of Economics
georgy.lukyanov@tse-fr.eu

Abstract

We study sequential social learning with endogenous information acquisition when agents have a taste for nonconformity. Each agent observes predecessors' actions, decides whether to acquire a private signal (and how precise it should be), and then chooses between two actions. Payoffs value correctness and include a bonus for taking the less popular action among predecessors; because this bonus depends only on observed popularity, the equilibrium analysis avoids fixed points in anticipated popularity and preserves standard Bayesian updating. In a Gaussian–quadratic setting, optimal actions follow posterior thresholds that tilt against the majority, and we solve the precision choice problem. Whenever the no-signal decision aligns with the observed majority, stronger contrarian motives weakly raise the value of information and expand the set of histories in which agents invest. We provide compact comparative statics for thresholds, action probabilities, and the precision argmax, a local welfare-and-information treatment, and applications to scientific priority races, cultural diffusion, and online platforms.

Keywords: social learning; information cascades; endogenous information acquisition; nonconformity; popularity; Bayesian thresholds.

JEL codes: D83; C72; D82; D85.

1 Introduction

Sequential social learning shows how rational agents can end up herding on the basis of observed actions. Once the information contained in past actions dominates what a typical private signal would add, later actions often stop reacting to new information. Yet in many settings agents have a taste for doing something different from the crowd: scientists may seek distinctiveness, consumers may value uniqueness, and political actors may position themselves against a majority. This paper studies how such contrarian motives interact with social learning when information acquisition is itself a choice.

We develop a simple, fully specified model in which each agent arrives sequentially, decides whether to buy a private signal and how precise it should be, and then chooses a binary action.

*We are grateful to Alexis Belianin, Emiliano Catonini, Darina Cheredina, Markus Gebauer, Olivier Gossner, Vitalijs Jascisens, Margarita Kirneva, Yukio Koriyama, Fabian Slonimczyk, and Alexey Verenikin for helpful comments and discussions. We also thank the participants of the ICEF research seminar for valuable feedback. Financial support from the French National Research Agency (ANR) under the "Investissements d'Avenir" program (LabEx Ecodec, ANR-11-LABX-0047) is gratefully acknowledged. All remaining errors are our own.

Agents value correctness and also derive a bonus from taking the less popular action, where popularity is the fraction of predecessors who chose a given action. The bonus is proportion-based, increasing in the size of the opposing crowd, and depends only on observed past actions. This design avoids any fixed-point in "anticipated popularity" and keeps Bayesian inference disciplined by the actual history. We work in a Gaussian–quadratic specification for tractability and transparency.

The analysis proceeds in two steps. First, we characterize action choice. Optimal behavior takes the form of a posterior (equivalently, log-likelihood ratio) threshold that depends on the observed popularity of the action (Lemma 4.1). The indifference cutoff shifts linearly with popularity and with the intensity of contrarian preferences, tilting decisions away from the observed majority (Proposition 4.2). We also quantify how the associated signal threshold moves with these parameters and how the induced action probabilities respond (Corollary 5.1 and Corollary 5.2).

Second, we study information acquisition. We formulate the precision choice problem with a fixed entry cost and a convex cost of precision and establish existence of an optimal precision together with an interior first-order condition in the Gaussian case (Proposition 4.3). We then define the net value of information at a given history and show that, whenever the no-signal action coincides with the observed majority, this value is (weakly) increasing in the strength of contrarian motives; hence the investment region expands with contrarian intensity in that empirically relevant region (Proposition 4.4). We also provide regularity results for the precision argmax correspondence—upper hemicontinuity in primitives, continuity of the value function, and local continuity (and, under standard conditions, differentiability) of the interior optimizer (Proposition 5.3).

Taken together, these results deliver a clear positive message. Contrarian motives shift decision thresholds against the crowd and thereby sustain experimentation precisely in those histories where, absent contrarianism, agents would be more prone to follow predecessors. At the same time, by making majority actions harder to choose, stronger contrarian motives can reduce the incidence of majority choices even when the observed crowd is informative.

2 Related literature

This paper connects three strands of work: social learning from actions, endogenous information acquisition, and preference-based (non)conformity.

The baseline we build on is the sequential social-learning tradition in which actions reveal private signals and can trigger herding and cascades. Foundational contributions include Banerjee (1992) and Bikhchandani et al. (1992), with sharper characterizations in Smith and Sørensen (2000, 2011). Extensions examine learning with partial observability and networks, e.g. Acemoglu et al. (2011) and Dasaratha and He (2020). A complementary literature studies why learning fails under misspecification or coarse inference even without payoff externalities, including Bohren (2016), Frick et al. (2020), Eyster et al. (2014), Bohren et al. (2019), and Kartik and Van Weelden (2020). Our contribution keeps Bayesian inference fully disciplined by observed actions and asks how a taste for nonconformity shifts the thresholds that govern action choice and, in turn, the informativeness of actions and the incentives to gather information.

We adopt a standard "fixed + convex" information technology to endogenize whether agents buy signals and how precise these signals are. Classic work on timing and informational externalities in sequential environments includes Chamley and Gale (1994). For formal treatments of information choice and precision, see Vives (2008) and Veldkamp (2011); recent empirical evidence on belief-dependent information acquisition in markets is provided by Lu et al. (2019). Within this tradition, our novelty is to show how observable, taste-based contrarian motives tilt decision thresholds and thereby expand the region of beliefs at which information acquisition is worthwhile, while preserving closed-form characterizations of action thresholds and regularity of the precision problem.

Classic conformity arises when actions signal type or status, as in Bernheim (1994), or when identity considerations affect payoffs, as in Akerlof and Kranton (2000) and Bénabou and Tirole (2011). We invert the sign and study *contrarian* preferences within a sequential learning game. Related models of distinctiveness and anti-conformism outside the canonical cascade environment include Golman et al. (2017) and Touboul (2019). A separate set of papers blends social learning with technological payoff externalities (coordination/complementarities), e.g. Dasgupta et al. (2011) and Ali and Kartik (2012). Our approach is different:[1] we add a taste-based bonus tied purely to *observed* popularity among predecessors, so the inference problem and equilibrium mapping from histories to beliefs remain standard, and there is no fixed-point in anticipated popularity. This modeling choice isolates how minority-seeking shifts cutoffs and affects information incentives without conflating preference shifts with payoff externalities.

Finally, it is useful to clarify when motives are contrarian rather than conformist. In applications where distinctiveness itself carries rents or identity value—scientific priority races, niche political positioning, or consumption domains emphasizing uniqueness—the minority action is privately valuable even holding informational incentives fixed. Our specification captures exactly this case by rewarding the observed minority choice, while leaving the Bayesian learning channel unchanged.

The paper is organized as follows. Section 3 presents the environment and assumptions, including the popularity-based contrarian preference and the information technology. Section 4 develops the equilibrium analysis: threshold behavior for actions, the precision choice problem with existence and interior FOC. Section 5 analyzes comparative statics. Section 6 provides a local welfare-and-information treatment, including the evaluator's value of information and an action informativeness measure. Section 7 discusses applications and motivation in science, culture, and online platforms. Section 8 concludes. Appendices collect full proofs and robustness notes.

3 Model

We study a canonical sequential social-learning environment in which agents endogenously choose whether to acquire information (and how precise it is) before taking a binary action. Agents value *correctness* and, in addition, have a taste for *nonconformity* that depends on the observed popularity of the action they choose among predecessors.

[1] In coordination/complementarity settings (e.g., Ali and Kartik (2012)), an agent's payoff depends on others' *actions*. Here, preferences add a taste term that depends only on the *observed popularity* of one's own chosen action, leaving Bayesian inference and the history-to-belief mapping unchanged.

3.1 Environment and information

Time is discrete, with agents $t = 1, 2, \ldots$ arriving sequentially. The state $\theta \in \{0, 1\}$ is fixed and drawn once at $t = 1$. The common prior is $\mu_1 \equiv \Pr(\theta = 1) = \frac{1}{2}$ (all results extend to any $\mu_1 \in (0, 1)$ with only notational changes). Each agent observes the entire history of past actions $h_{t-1} = (a_1, \ldots, a_{t-1})$, where $a_i \in \{0, 1\}$.

Before acting, agent t may *acquire a private signal* of endogenous precision $\rho_t \geq 0$. If $\rho_t = 0$, no signal is acquired. If $\rho_t > 0$, the signal is Gaussian:[2]

$$s_t \mid \theta \;\sim\; \mathcal{N}\big(\theta,\; 1/\rho_t\big),$$

conditionally independent across agents given θ. Let $\ell_t(s_t; \rho_t)$ denote the log-likelihood ratio (LLR), which in this Gaussian specification equals $\ell_t(s_t; \rho_t) = \rho_t\,(s_t - \frac{1}{2})$ (equal variances and means 0 and 1).

Acquiring any signal entails a fixed cost $F > 0$. Precision is chosen at a convex cost $C(\rho) = \frac{c}{2}\rho^2$ with $c > 0$.[3] Thus the total information cost is $F\,\mathbf{1}\{\rho > 0\} + \frac{c}{2}\rho^2$.

3.2 Preferences and contrarian utility

An agent's payoff has two components:

- *Correctness:* the agent obtains 1 if $a_t = \theta$ and 0 otherwise.
- *Nonconformity:* let $N_{t-1}(a) = \sum_{i=1}^{t-1} \mathbf{1}\{a_i = a\}$ and define the observed popularity of action a among predecessors at time t by

$$p_t(a) \;=\; \begin{cases} \frac{N_{t-1}(a)}{t-1}, & t \geq 2, \\ \frac{1}{2}, & t = 1, \end{cases}$$

  so that $p_t(0) + p_t(1) = 1$ for $t \geq 2$ and both actions are treated symmetrically at $t = 1$. The contrarian bonus is proportional to how *unpopular* the chosen action is:[4]

$$b\big(p_t(a_t)\big) \;=\; k\,\big(1 - p_t(a_t)\big), \qquad k \geq 0.$$

Hence agent t's (ex post) utility is

$$u_t(a_t, \theta; h_{t-1}) \;=\; \mathbf{1}\{a_t = \theta\} \;+\; k\,\big(1 - p_t(a_t)\big).$$

[2]Normality is chosen for closed-form thresholds and clean comparative statics. All results that rely on single-crossing/MLRP extend to any signal family with MLRP; formulas then become implicit but the qualitative monotone effects are unchanged.

[3]Any strictly convex C with $C(0) = C'(0) = 0$ yields the same comparative statics; quadratic cost is a transparent benchmark. The fixed cost $F > 0$ generates the empirically relevant corner at $\rho = 0$; setting $F = 0$ collapses to a pure convex case with the same Euler condition.

[4]Linear $b(p) = k(1-p)$ is a convenient normalization. Any smooth, strictly decreasing b with $b'(p) < 0$ delivers the same first-order threshold tilt around interior beliefs; linearity makes the dependence on p_t and k fully transparent.

The bonus depends only on *observed predecessors' actions*. It does not depend on contemporaneous or future actions, so there is no fixed-point in defining the bonus.[5]

3.3 Timing

At date t:

1. The public history h_{t-1} is observed. The public belief is $\mu_t = \Pr(\theta = 1 \mid h_{t-1})$ (defined below).

2. Agent t chooses a precision $\rho_t \geq 0$, pays $F\,\mathbf{1}\{\rho_t > 0\} + \frac{c}{2}\rho_t^2$, and if $\rho_t > 0$ observes $s_t \sim \mathcal{N}(\theta, 1/\rho_t)$.

3. Agent t chooses $a_t \in \{0,1\}$ and receives $u_t(a_t, \theta; h_{t-1})$.

4. a_t is publicly revealed; the history updates to $h_t = (h_{t-1}, a_t)$.

3.4 Beliefs and posteriors

Let σ denote a strategy profile (defined below). Given σ, the public belief after history h_{t-1} is

$$\mu_t \;=\; \Pr(\theta = 1 \mid h_{t-1}; \sigma).$$

If agent t acquires precision ρ_t and observes s_t, her posterior is

$$\Pr(\theta = 1 \mid h_{t-1}, s_t; \sigma) \;=\; \Lambda\!\left(\log \frac{\mu_t}{1-\mu_t} \;+\; \ell_t(s_t; \rho_t)\right),$$

where $\Lambda(x) = \frac{e^x}{1+e^x}$ is the logistic function and $\ell_t(s_t; \rho_t) = \rho_t\,(s_t - \frac{1}{2})$ is the Gaussian LLR. If $\rho_t = 0$ (no signal) the posterior equals the public belief, i.e., $\Pr(\theta = 1 \mid h_{t-1}; \sigma) = \mu_t$.

3.5 Strategies and equilibrium

A (pure) strategy for agent t is a pair of measurable maps

$$\big(\varphi_t,\ \alpha_t\big) : \ \mathcal{H}_{t-1} \to \mathbb{R}_+ \times \mathcal{A},$$

where $\mathcal{H}_{t-1}$ is the set of feasible histories and $\mathcal{A}$ is the set of Borel-measurable action rules $\alpha_t : \mathbb{R} \cup \{\varnothing\} \to \{0,1\}$ mapping either a realized signal $s_t \in \mathbb{R}$ (if $\rho_t = \varphi_t(h_{t-1}) > 0$) or the null input $\varnothing$ (if $\rho_t = 0$) into an action. We write $\rho_t = \varphi_t(h_{t-1})$.

An *assessment* (σ, μ) consists of a profile of strategies $\sigma = \{\varphi_t, \alpha_t\}_{t\geq 1}$ and a system of public beliefs $\{\mu_t\}_{t\geq 1}$.

[5] If the bonus depended on *expected* contemporaneous popularity, equilibrium would embed a fixed-point in popularity beliefs and could generate multiplicity unrelated to learning. Tying $b(\cdot)$ to *observed predecessors* isolates the informational channel.

Equilibrium concept. An assessment (σ, μ) is a *Perfect Bayesian Equilibrium*[6] *(PBE)* if: (i) For each t and history h_{t-1}, $\big(\varphi_t(h_{t-1}), \alpha_t(\cdot \mid h_{t-1})\big)$ maximizes agent t's expected utility net of information costs, given (σ, μ); (ii) μ is obtained from Bayes' rule whenever applicable, given σ and the model's likelihoods; off-path beliefs are specified in a way consistent with standard refinements. We adopt the tie-breaking convention that an indifferent agent picks the action with the higher contrarian bonus (the less popular action);[7] results are unchanged under any fixed tie-breaking rule.

3.6 Popularity statistic and basic definitions

Let $p_t(a)$ be as defined above. We use $p_t \equiv p_t(1)$ as a shorthand for the popularity of action 1 at time t, so $p_t(0) = 1 - p_t$. The popularity statistic is fully observable at time t.

Definition 3.1. At history h_{t-1}, agent t's action is *informative* if, conditional on the public belief μ_t and the strategy profile σ, there exist signal realizations leading to each action with strictly positive probability; equivalently, both $\Pr(a_t = 1 \mid \theta = 1, h_{t-1}; \sigma)$ and $\Pr(a_t = 1 \mid \theta = 0, h_{t-1}; \sigma)$ lie strictly between 0 and 1.

Definition 3.2. A 1-*cascade* occurs at history h_{t-1} if $a_t = 1$ almost surely under σ for both states; a 0-*cascade* is defined analogously. In a cascade, actions are uninformative.[8]

3.7 Standing assumptions

We maintain the following assumptions throughout the paper:

(A1) **Signals.** Conditional on θ, private signals are conditionally independent across agents; if precision $\rho > 0$ is chosen, $s \mid \theta \sim \mathcal{N}(\theta, 1/\rho)$.

(A2) **Costs.** $F > 0$ and $c > 0$; the cost of precision is $C(\rho) = \frac{c}{2}\rho^2$.

(A3) **Preferences.** The contrarian bonus is $b(p) = k(1 - p)$ with $k \geq 0$, applied to the observed popularity of the *chosen* action among predecessors.

(A4) **Observability and common knowledge.** The entire action history h_{t-1} is public; signals are private; (F, c, k) and the model are common knowledge.

These assumptions, combined with the Gaussian specification, imply that posterior beliefs are monotone in s_t (MLRP)[9] and that optimal action rules can be represented as posterior- or LLR-thresholds in equilibrium. Formal statements and proofs appear in Section 4.

[6] All statements carry over to sequential equilibrium. We restrict to pure strategies throughout; randomization matters only at indifference points.

[7] Any fixed tie-breaking rule yields the same results generically; it only bites on knife-edge histories (measure zero under continuous signals).

[8] With endogenous information, contrarian motives can disrupt would-be cascades by restoring incentives to acquire signals. Proposition 4.4 formalizes the region where this happens.

[9] With asymmetric priors $\mu_1 \neq \frac{1}{2}$ all formulas go through with the log-odds shift $L_t = \log \frac{\mu_t}{1-\mu_t}$; we normalize $\mu_1 = \frac{1}{2}$ for symmetry in the exposition.

4 Equilibrium analysis

We characterize optimal action choice as a posterior/LLR threshold, derive the dependence of the threshold on the observed popularity p_t, and analyze the information acquisition problem (existence, first-order condition, and comparative statics in k for the value of information under clear baseline conditions).

Throughout the section fix a date t and a public history h_{t-1} with associated public belief $\mu_t \in (0,1)$, log-odds $L_t := \log \frac{\mu_t}{1-\mu_t}$, and observed popularity $p_t(1) = N_{t-1}(1)/(t-1)$ (with the convention $p_1(1) = 1/2$). Write $p_t := p_t(1)$ and $p_t(0) = 1 - p_t$.

4.1 Action choice: posterior and LLR thresholds

At history h_{t-1}, the bonus differential from choosing 1 rather than 0 is

$$\Delta_t \;:=\; b\big(p_t(1)\big) - b\big(p_t(0)\big) \;=\; k\,(1-p_t) - k\,p_t \;=\; k\,(1-2p_t).$$

Let c_t be the posterior cutoff that makes the agent indifferent between actions.

Lemma 4.1. *Fix a history h_{t-1}. For any chosen precision $\rho \geq 0$ and any realized signal s, the difference in expected utility between actions 1 and 0 equals*

$$\underbrace{2\,\Pr(\theta = 1 \mid h_{t-1}, s) - 1}_{\textit{correctness term}} + \underbrace{\Delta_t}_{\textit{contrarian bonus difference}}.$$

Hence the optimal action is a posterior threshold: choose $a_t = 1$ iff

$$\Pr(\theta = 1 \mid h_{t-1}, s) \;\geq\; c_t, \qquad c_t \;=\; \frac{1-\Delta_t}{2} \;=\; \frac{1}{2} + k\Big(p_t - \frac{1}{2}\Big).$$

Equivalently, in LLR units choose $a_t = 1$ iff

$$L_t \;+\; \ell_t(s;\rho) \;\geq\; \tau_t, \qquad \tau_t \;:=\; \log\frac{c_t}{1-c_t}.$$

In the Gaussian case with $\ell_t(s;\rho) = \rho\,(s - \frac{1}{2})$, this is the signal threshold rule $a_t = 1$ iff $s \geq s_t^\star(\rho)$ where

$$s_t^\star(\rho) \;=\; \frac{1}{2} \;+\; \frac{\tau_t - L_t}{\rho} \qquad \text{(with the convention that } s_t^\star(0) = +\infty \text{ if } \mu_t < c_t,\ -\infty \text{ if } \mu_t > c_t\text{)}.$$

Proof. Let $\pi(s) := \Pr(\theta = 1 \mid h_{t-1}, s)$ denote the posterior. The expected correctness payoff of $a = 1$ is $\pi(s)$ and of $a = 0$ is $1 - \pi(s)$, so the difference is $2\pi(s) - 1$. The expected contrarian bonus of $a = 1$ is $k(1 - p_t)$ and of $a = 0$ is $k(1 - p_t(0)) = k\,p_t$, so the bonus difference is $\Delta_t = k(1 - 2p_t)$, which does not depend on s or ρ. Therefore the agent prefers 1 to 0 iff $2\pi(s) - 1 + \Delta_t \geq 0$, i.e., $\pi(s) \geq (1 - \Delta_t)/2 =: c_t$. Since posteriors in a binary-state, MLRP environment are strictly increasing in the LLR, the posterior cutoff implies the LLR cutoff $L_t + \ell_t(s;\rho) \geq \tau_t := \log \frac{c_t}{1-c_t}$. In the Gaussian case $\ell_t(s;\rho) = \rho(s - \frac{1}{2})$ gives the claimed s-threshold. □

Proposition 4.2. *For any history, $c_t = \frac{1}{2} + k(p_t - \frac{1}{2})$. Hence: (i) c_t is strictly increasing in p_t when $k > 0$; (ii) if $p_t > \frac{1}{2}$ then $c_t > \frac{1}{2}$ and c_t is strictly increasing in k; if $p_t < \frac{1}{2}$ then $c_t < \frac{1}{2}$ and c_t is strictly decreasing in k; (iii) if $p_t = \frac{1}{2}$, then $c_t = \frac{1}{2}$ for all k.*

Proof. Immediate from Lemma 4.1 since $\Delta_t = k(1 - 2p_t)$. □

4.2 Information acquisition: existence, FOC, and probabilities

Given (μ_t, p_t) and $\rho \geq 0$, define the *gross expected payoff*

$$G_t(\mu_t, \rho; p_t, k) := \underbrace{\mu_t \; \Pr(a_t = 1 \mid \theta = 1) + (1 - \mu_t) \; \Pr(a_t = 0 \mid \theta = 0)}_{\text{correctness}} + k \; \underbrace{\mathbb{E}[1 - p_t(a_t)]}_{\text{contrarian bonus}},$$

where probabilities are induced by the Gaussian signal and the threshold rule of Lemma 4.1. The agent then solves

$$\max_{\rho \geq 0} \; \Big\{ G_t(\mu_t, \rho; p_t, k) \; - \; \frac{c}{2}\rho^2 \; - \; F \, \mathbf{1}\{\rho > 0\} \Big\}.$$

We first make the Gaussian probabilities explicit. Let

$$s_t^\star(\rho) = \frac{1}{2} + \frac{\tau_t - L_t}{\rho}, \qquad z_0(\rho) := s_t^\star(\rho)\sqrt{\rho}, \qquad z_1(\rho) := \big(s_t^\star(\rho) - 1\big)\sqrt{\rho},$$

and let Φ and ϕ denote the standard normal cdf and pdf, respectively. Then

$$\begin{aligned} \Pr(a_t = 1 \mid \theta = 1) &= 1 - \Phi\big(z_1(\rho)\big), \\ \Pr(a_t = 1 \mid \theta = 0) &= 1 - \Phi\big(z_0(\rho)\big), \\ \Pr(a_t = 0 \mid \theta = 0) &= \Phi\big(z_0(\rho)\big), \end{aligned}$$

and the unconditional choice probability is

$$\Pr(a_t = 1) = \mu_t \, [1 - \Phi(z_1(\rho))] + (1 - \mu_t) \, [1 - \Phi(z_0(\rho))].$$

Hence the expected contrarian bonus equals

$$\mathbb{E}[1 - p_t(a_t)] \; = \; (1 - p_t) \; \Pr(a_t = 1) \; + \; p_t \, \big(1 - \Pr(a_t = 1)\big) \; = \; p_t \; + \; (1 - 2p_t) \; \Pr(a_t = 1).$$

Proposition 4.3. *For any (μ_t, p_t, k), the information acquisition problem admits a solution $\rho_t^\star \geq 0$.*[10] *Moreover, if $\rho_t^\star > 0$ solves the problem, then it satisfies the first-order condition*

$$\frac{\partial}{\partial \rho} G_t(\mu_t, \rho; p_t, k)\Big|_{\rho = \rho_t^\star} \; = \; c \, \rho_t^\star,$$

[10] Global concavity of the objective in ρ need not hold. Our comparative-statics results rely on the argmax correspondence (compactness/upper hemicontinuity) rather than on single-peakedness of $J(\rho; \cdot)$.

where

$$\frac{\partial}{\partial\rho}\,G_t(\mu_t,\rho;p_t,k) = \overbrace{\mu_t\,\phi\big(z_1(\rho)\big)\,(-z_1'(\rho)) \;+\; (1-\mu_t)\,\phi\big(z_0(\rho)\big)\,z_0'(\rho)}^{\text{marginal correctness gain}}$$
$$+\;\; k\,(1-2p_t)\;\underbrace{\Big[\mu_t\,\phi\big(z_1(\rho)\big)\,(-z_1'(\rho)) \;+\; (1-\mu_t)\,\phi\big(z_0(\rho)\big)\,(-z_0'(\rho))\Big]}_{\frac{\partial}{\partial\rho}\Pr(a_t=1)}.$$

Here $z_0'(\rho)$ and $z_1'(\rho)$ are given by

$$z_0'(\rho) = \frac{s_t^\star(\rho)}{2\sqrt{\rho}} + \sqrt{\rho}\,s_t^{\star\prime}(\rho), \qquad z_1'(\rho) = \frac{s_t^\star(\rho)-1}{2\sqrt{\rho}} + \sqrt{\rho}\,s_t^{\star\prime}(\rho), \quad s_t^{\star\prime}(\rho) = -\frac{\tau_t - L_t}{\rho^2}.$$

Proof sketch. Continuity and the quadratic cost imply existence of a maximizer. In the Gaussian case, write the state-conditional choice probabilities using the signal threshold $s_t^\star(\rho)$. Differentiating G_t yields the Euler condition $\partial_\rho G_t = c\rho$ at any interior optimum. Full derivatives appear in Appendix A. □

4.3 Value of information and its monotonicity in k

Define the (net) value of information at (μ_t, p_t, k) by

$$\Phi_t(\mu_t,p_t;k) \;:=\; \max_{\rho\ge 0}\Big\{G_t(\mu_t,\rho;p_t,k) - \frac{c}{2}\rho^2\Big\} \;-\; G_t(\mu_t,0;p_t,k).$$

Observe that $G_t(\mu_t,\rho;p_t,k) = \mathrm{Corr}_t(\mu_t,\rho) \;+\; k\,M_t(\mu_t,\rho;p_t)$, where $M_t(\mu_t,\rho;p_t) := \mathbb{E}[1-p_t(a_t)]$. Hence

$$\Phi_t(\mu_t,p_t;k) = \sup_{\rho\ge 0}\Big\{A_t(\rho) \;+\; k\,B_t(\rho) \;-\; \frac{c}{2}\rho^2\Big\},$$
$$A_t(\rho) := \mathrm{Corr}_t(\mu_t,\rho) - \mathrm{Corr}_t(\mu_t,0),$$
$$B_t(\rho) := M_t(\mu_t,\rho;p_t) - M_t(\mu_t,0;p_t).$$

The sign of $B_t(\rho)$ depends on whether the *no-signal* action aligns with the observed majority or minority. The following proposition isolates a robust region where Φ_t is (weakly) increasing in k and therefore the investment region expands with k.

Proposition 4.4. *Fix (μ_t, p_t) with $p_t \neq \frac{1}{2}$ and let a_t^0 be the action chosen when $\rho = 0$ (i.e., $a_t^0 = 1$ iff $\mu_t \ge c_t$). If a_t^0 coincides with the observed* majority *action at h_{t-1} (i.e., $a_t^0 = 1$ when $p_t > \frac{1}{2}$, or $a_t^0 = 0$ when $p_t < \frac{1}{2}$), then for all $\rho \ge 0$, $B_t(\rho) \ge 0$ and thus $k \mapsto \Phi_t(\mu_t,p_t;k)$ is (weakly) increasing. Consequently, for any fixed $F > 0$, the set $\{\ k \ge 0 :\ \Phi_t(\mu_t,p_t;k) > F\ \}$ is an interval of the form $[k_0,\infty)$ (possibly $k_0 = 0$ or $+\infty$).*

Proof sketch. When the baseline action follows the observed majority, raising ρ weakly *reduces* the expected bonus shortfall at $\rho = 0$, so G_t has increasing differences in (ρ, k). Taking the supremum over $\rho \ge 0$ preserves monotonicity in k. See Appendix A for details. □

Remark 4.5. If the *no-signal* action coincides with the *minority* (e.g., $p_t > \frac{1}{2}$ and $\mu_t < c_t$, so $a_t^0 = 0$), then $P_1(0) = 0$ and the same algebra yields $B_t(\rho) \le 0$ for all $\rho \ge 0$. In that case $k \mapsto \Phi_t(\mu_t, p_t; k)$ need not be increasing. Proposition 4.4 cleanly separates the region where monotonic expansion in k is guaranteed from the complementary region where no general monotonicity can be claimed.

5 Comparative statics

We collect a few immediate implications for thresholds, signal cutoffs, and the precision argmax correspondence.

Corollary 5.1. *For any history with popularity $p_t \in [0,1]$ and $k \ge 0$:*

$$\frac{\partial c_t}{\partial k} = p_t - \frac{1}{2}, \qquad \frac{\partial \tau_t}{\partial k} = \frac{p_t - \frac{1}{2}}{c_t(1-c_t)}, \qquad \frac{\partial s_t^\star}{\partial k} = \frac{p_t - \frac{1}{2}}{\rho\, c_t(1-c_t)}.$$

Hence, if $p_t > \frac{1}{2}$ (action 1 is the observed majority), then c_t and τ_t increase in k and the signal threshold $s_t^\star$ shifts up *in k (it becomes harder to choose the majority action 1). If $p_t < \frac{1}{2}$, the inequalities reverse. When $p_t = \frac{1}{2}$, all three derivatives are zero.*

Proof. From Lemma 4.1, $c_t = \frac{1}{2} + k(p_t - \frac{1}{2})$ so $\partial c_t/\partial k = p_t - \frac{1}{2}$. Since $\tau_t = \log \frac{c_t}{1-c_t}$, the chain rule gives $\partial \tau_t/\partial k = (\partial \tau_t/\partial c_t)(\partial c_t/\partial k) = \frac{1}{c_t(1-c_t)}(p_t - \frac{1}{2})$. Finally, $s_t^\star(\rho) = \frac{1}{2} + (\tau_t - L_t)/\rho$ implies $\partial s_t^\star/\partial k = (1/\rho)\, \partial \tau_t/\partial k$. □

Corollary 5.2. *Fix (μ_t, p_t, k, ρ) and let $P_1^\theta(\rho) := \Pr(a_t = 1 \mid \theta)$ under the Gaussian threshold. Then*

$$\frac{\partial}{\partial k} P_1^1(\rho) = -\,\phi\big(z_1(\rho)\big)\,\frac{1}{\sqrt{\rho}}\,\frac{p_t - \frac{1}{2}}{c_t(1-c_t)}, \qquad \frac{\partial}{\partial k} P_1^0(\rho) = -\,\phi\big(z_0(\rho)\big)\,\frac{1}{\sqrt{\rho}}\,\frac{p_t - \frac{1}{2}}{c_t(1-c_t)}.$$

Thus if $p_t > \frac{1}{2}$ (majority is action 1), increasing k reduces *the likelihood of choosing 1 in both states, and vice versa when $p_t < \frac{1}{2}$.*

Proof. $P_1^1(\rho) = 1 - \Phi(z_1(\rho))$ and $P_1^0(\rho) = 1 - \Phi(z_0(\rho))$ with $z_\iota(\rho) = \big(s_t^\star(\rho) - \iota\big)\sqrt{\rho}$, $\iota \in \{0,1\}$. By the chain rule, $\partial_k[1 - \Phi(z_\iota)] = -\phi(z_\iota)\, \partial_k z_\iota$ and $\partial_k z_\iota = \sqrt{\rho}\, \partial_k s_t^\star$. Substitute $\partial_k s_t^\star$ from Corollary 5.1. □

Proposition 5.3. *Let $x := (\mu_t, p_t, k)$ and define the objective*

$$J(\rho; x) \;:=\; G_t(\mu_t, \rho; p_t, k) \;-\; \frac{c}{2}\rho^2 \;-\; F\,\mathbf{1}\{\rho > 0\}, \qquad \rho \in [0, \infty).$$

(a) For each x, the argmax set $\Gamma(x) := \arg\max_{\rho \ge 0} J(\rho; x)$ is nonempty and compact.
(b) The correspondence $x \mapsto \Gamma(x)$ is upper hemicontinuous. In particular, the value function *$V(x) := \max_{\rho \ge 0} J(\rho; x)$ is continuous in x.*
(c) If for some x the maximizer is unique and interior ($\Gamma(x) = \{\rho^\star > 0\}$) and the strict second-order condition *holds, $\partial_{\rho\rho} G_t(\mu_t, \rho^\star; p_t, k) - c < 0$, then $\rho^\star$ depends continuously on x in a neighborhood of x. If, in addition, the FOC $\partial_\rho G_t(\mu_t, \rho^\star; p_t, k) = c\rho^\star$ has nonzero derivative in ρ at $\rho^\star$, then $\rho^\star$ is continuously differentiable in x by the implicit function theorem.*

Proof sketch. Compact truncation, upper semicontinuity at $\rho = 0$, and continuity elsewhere allow an application of Berge's maximum theorem; the implicit function theorem gives local smoothness under uniqueness and a strict SOC. Full details are in Appendix A. □

Remark 5.4. In general, $\rho_t^\star$ need not be monotone in k because the term $k(1-2p_t)\Pr(a_t = 1)$ can generate either increasing or decreasing differences in (ρ, k), depending on p_t. Proposition 4.4 shows that the *maximized value of information* Φ_t is (weakly) increasing in k whenever the no-signal action follows the observed majority; this does not, by itself, pin down the sign of $\mathrm{d}\rho_t^\star/\mathrm{d}k$.

6 Welfare and information content

We provide a local welfare accounting at a given public history h_{t-1} with state-belief $\mu_t \in (0,1)$ and observed popularity $p_t \in [0,1]$. We study (i) a myopic evaluator who places weight $\lambda \in [0,1]$ on the contrarian bonus (with $\lambda = 0$ corresponding to correctness-only evaluation and $\lambda = 1$ coinciding with the agent's taste), and (ii) the *information content* of the public action as a primitive measure of its usefulness for future learning.[11] Throughout, strategies and beliefs are as in Section 4.

6.1 Evaluator welfare and value-of-information accounting

Fix $\lambda \in [0,1]$. At (μ_t, p_t) and precision choice $\rho \geq 0$, define the evaluator's *one-step expected payoff*

$$W_t^\lambda(\mu_t, \rho; p_t, k) \;:=\; \overbrace{\mu_t \Pr(a_t = 1 \mid \theta = 1) + (1-\mu_t)\Pr(a_t = 0 \mid \theta = 0)}^{\text{correctness}} \\ +\; \lambda k \underbrace{\mathbb{E}[1 - p_t(a_t)]}_{\text{distinctiveness}} - \frac{c}{2}\rho^2 - F\,\mathbf{1}\{\rho > 0\},$$

where probabilities are induced by the Gaussian threshold in Lemma 4.1. The associated (myopic) evaluator *value of information* is

$$\Phi_t^\lambda(\mu_t, p_t; k) \;:=\; \max_{\rho \geq 0}\Big\{\mu_t \Pr(a_t = 1 \mid \theta = 1) + (1-\mu_t)\Pr(a_t = 0 \mid \theta = 0) + \lambda k\,\mathbb{E}[1 - p_t(a_t)] - \tfrac{c}{2}\rho^2\Big\}.$$

By construction, Φ_t^λ treats F as an entry cost (it cancels in the difference).

Proposition 6.1. *Fix (μ_t, p_t) with $p_t \neq \frac{1}{2}$ and let a_t^0 be the no-signal action (i.e., $a_t^0 = 1$ iff $\mu_t \geq c_t$ with $c_t = \frac{1}{2} + k(p_t - \frac{1}{2})$). Suppose a_t^0 coincides with the observed* majority *(i.e., $a_t^0 = 1$ when $p_t > \frac{1}{2}$, or $a_t^0 = 0$ when $p_t < \frac{1}{2}$). Then for every $\lambda \in [0,1]$, $k \mapsto \Phi_t^\lambda(\mu_t, p_t; k)$ is (weakly) increasing. Consequently, for any fixed $F > 0$ the myopic evaluator's investment region $\{k : \Phi_t^\lambda(\mu_t, p_t; k) > F\}$ is an interval $[k_\lambda, \infty)$ (possibly degenerate).*

Proof. The proof repeats the argument of Proposition 4.4 with k replaced by λk in the evaluator's objective. Write $P_1(\rho) := \Pr(a_t = 1)$ and recall $\mathbb{E}[1 - p_t(a_t)] = p_t + (1-2p_t)P_1(\rho)$. When $p_t > \frac{1}{2}$

[11] We deliberately avoid global welfare rankings across k. Aggregating over the endogenous belief path would require dynamic weights and preference aggregation beyond our scope; the local evaluator view and an action-informativeness metric are transparent and policy-relevant.

and $a_t^0 = 1$, we have $P_1(0) = 1$ and $(1 - 2p_t) < 0$, so for all $\rho \geq 0$,

$$\left[\mathbb{E}[1 - p_t(a_t)] - \mathbb{E}[1 - p_t(a_t)]\right]_{\rho=0} = (1 - 2p_t)\left(P_1(\rho) - 1\right) \geq 0.$$

Thus the map $\rho \mapsto$ correctness $+ \lambda k\,\mathbb{E}[1 - p_t(a_t)] - \frac{c}{2}\rho^2$ has (weakly) increasing differences in (ρ, k) for any fixed $\lambda \in [0, 1]$, and so its supremum over $\rho \geq 0$ is (weakly) increasing in k. The case $p_t < \frac{1}{2}$ is symmetric. Continuity and the interval statement follow by Berge's maximum theorem. □

6.2 Information content of the public action

We quantify the informativeness of the period-t action for the state by the total variation distance[12] between the state-conditional action distributions:

$$\mathcal{I}_t(\rho) \;:=\; \left|\Pr(a_t = 1 \mid \theta = 1) - \Pr(a_t = 1 \mid \theta = 0)\right|.$$

This measure equals zero iff the action is uninformative about θ.

Lemma 6.2. *In the Gaussian environment with the threshold decision rule of Lemma 4.1, $\mathcal{I}_t(\rho) = 0$ iff $\rho = 0$. For every $\rho > 0$, $\mathcal{I}_t(\rho) > 0$.*

Proof. For $\rho = 0$, the action is chosen from the prior alone and is (weakly) deterministic, so $\Pr(a_t = 1 \mid \theta = 1) = \Pr(a_t = 1 \mid \theta = 0)$ and $\mathcal{I}_t(0) = 0$. For $\rho > 0$, the signal threshold $s_t^\star(\rho) \in \mathbb{R}$ is finite. With $s \mid \theta \sim \mathcal{N}(\theta, 1/\rho)$, we have

$$\Pr(a_t = 1 \mid \theta = 1) = 1 - \Phi\big((s_t^\star(\rho) - 1)\sqrt{\rho}\big), \quad \Pr(a_t = 1 \mid \theta = 0) = 1 - \Phi\big(s_t^\star(\rho)\sqrt{\rho}\big).$$

Since $s_t^\star(\rho)\sqrt{\rho} - (s_t^\star(\rho) - 1)\sqrt{\rho} = \sqrt{\rho} > 0$, the monotonicity of Φ implies $\Pr(a_t = 1 \mid \theta = 1) > \Pr(a_t = 1 \mid \theta = 0)$, hence $\mathcal{I}_t(\rho) > 0$. □

The next proposition links information content to contrarian intensity through the investment margin highlighted earlier.

Proposition 6.3. *Fix (μ_t, p_t) satisfying the baseline-majority condition of Proposition 4.4. Let $\rho_t^\star(k)$ be an equilibrium precision choice at (μ_t, p_t) for contrarian intensity k. Then there exists $k_0 \in [0, \infty]$ such that $\rho_t^\star(k) = 0$ for all $k < k_0$ and $\rho_t^\star(k) > 0$ for all $k > k_0$. Moreover, $k \mapsto \mathcal{I}_t\big(\rho_t^\star(k)\big)$ is (weakly) increasing and has a (weakly) positive jump at k_0 whenever $k_0 < \infty$.*

Proof. By Proposition 4.4, $\Phi_t(\mu_t, p_t; k)$ is increasing in k and the agent's investment region $\{k : \Phi_t(\mu_t, p_t; k) > F\}$ is an interval $[k_0, \infty)$, so $\rho_t^\star(k) = 0$ for $k < k_0$ and $\rho_t^\star(k) > 0$ for $k > k_0$. The mapping $k \mapsto \mathcal{I}_t(\rho_t^\star(k))$ equals 0 on $[0, k_0)$ and is strictly positive on (k_0, ∞) by Lemma 6.2. Upper hemicontinuity of the argmax (Proposition 5.3) yields right-continuity of $\rho_t^\star(k)$ away from the entry point; at k_0 the function $\mathcal{I}_t(\rho_t^\star(k))$ has a (weakly) positive jump from 0 to $\mathcal{I}_t(\bar{\rho})$ for any interior maximizer $\bar{\rho} \in \Gamma(\mu_t, p_t, k_0^+)$. □

[12] Total variation equals the L^1 distance between the two Bernoulli action distributions and pins down the optimal one-shot use of a_t for inference. Other measures (e.g., mutual information) yield the same qualitative conclusions here.

Remark 6.4. Our welfare objects are *local* (one-step) at a given history. A full dynamic welfare ranking across different values of k would require aggregating over the endogenous path of public beliefs induced by equilibrium play, which lies beyond our scope. Propositions 6.1 and 6.3 show that, on the empirically relevant baseline-majority region, stronger contrarian motives weakly expand (i) the evaluator's myopic value of information for any $\lambda \in [0, 1]$ and (ii) the set of histories at which actions are strictly informative.

7 Applications and motivation

Scientific progress, cultural diffusion, and online attention markets all feature a tension between the pull of consensus and the private appeal of standing apart. Our model distills this tension into a transparent, history-based bonus for taking the less popular action and shows how that bonus shifts decision thresholds and information choices in ways that map directly into observable behaviors.

A scholar's evidence—from a new experiment, dataset, or argument—may clash with prevailing orthodoxy. The personal payoff to nonconformity can be substantial in such environments. As Max Planck wrote in his *Scientific Autobiography*:

> "A new scientific truth does not triumph by convincing its opponents and making them see the light, but rather because its opponents eventually die, and a new generation grows up that is familiar with it" (Planck, 1949, pp. 33–34).

Kuhn echoed the same theme in *The Structure of Scientific Revolutions*:

> "Paradigm change cannot be justified by proof. Instead it must be decided by persuasion and conversion, and often by generational replacement" (Kuhn, 1962, p. 151).

Our positive results formalize when such contrarian motives keep inquiry alive rather than derail standards of evidence. The action rule is a clean posterior (LLR) threshold that *tilts against the observed majority* (Lemma 4.1 and Proposition 4.2). Thus, in fields where a topic has recently become popular (p_t high), a stronger taste for distinctiveness (higher k) mechanically raises the cutoff for following that majority, making it more likely that dissenting findings are voiced rather than suppressed. On the information margin, whenever the no-signal decision coincides with the observed majority—arguably the empirically relevant case near dominant paradigms—increasing k *expands the investment region*: more teams choose to collect evidence (Proposition 4.4). Together, these effects predict bursts of experimentation exactly when orthodoxy is strong,[13] consistent with the idea that distinctiveness rents encourage outsiders to keep testing dominant views.

Two measurement implications follow. First, as a topic's popularity rises, the probability of majority actions should decline at the margin (Corollary 5.2), even holding public belief fixed. Second, entry into data collection (or running new experiments) should be more responsive to the observed crowd than to fine differences in priors, lining up with our region-wise monotonicity in k (Proposition 4.4). Regularity of the precision choice (§5, Proposition 5.3) ensures these comparative statics are robust to small shifts in primitives.

[13] Formally, when the no-signal action follows the majority, increasing k (salience of distinctiveness) weakly expands the investment region (Proposition 4.4); interior precision and action probabilities then move as in Corollaries 5.1–5.2.

In cultural markets, status often derives from distinctiveness: early adopters and "outsiders" reap private benefits from nonconformity. Our threshold characterization implies that visible popularity (p_t) crowds some agents *away* from the mainstream choice, generating cycling and diversity in observed actions even without technological externalities. When a genre, style, or meme is already popular, higher k makes additional adoptions less likely at the margin (Corollary 5.1 and Corollary 5.2); yet the same contrarian force invites more "tests" of alternatives, expanding exploration (Proposition 4.4). This yields a simple macro-pattern: mainstream surges coincide with more attempts at differentiation and more informative minority actions, whereas very strong contrarianism can push behavior into low-value deviation for its own sake.

On platforms where actions (likes, shares, reposts) are public, p_t is highly salient. Our model predicts two opposite forces as p_t grows: (i) a tilt away from boosting already-popular content (threshold shift), and (ii) a larger set of histories in which users invest in "search" (click-throughs, reading, creating competing posts) to find contrarian content (investment-region expansion). Design choices that make popularity more or less salient[14] therefore change both *what* gets amplified and *how much* users investigate alternatives. Because our results keep the inference channel Bayesian, these design predictions are separate from algorithmic feed-back effects: they arise even in simple, transparent feeds.

The results yield low-commitment instruments for sustaining informativeness without prescribing beliefs. In environments prone to premature herding, modest visibility of popularity (or small symbolic rewards for minority contributions) can replicate the threshold tilt and enlarge the range of histories in which agents test alternatives (Proposition 4.4). Conversely, when fragmentation is a concern, dampening the salience of p_t can mitigate excessive contrarianism. These are positive implications; we refrain from ranking k globally.

Our applications rely on the taste-based bonus being tied to *observed* popularity among predecessors. This choice avoids fixed-point complications in expectations about popularity and keeps inference disciplined by the actual history. It is precisely this clarity that lets us connect quotes like Planck's and Kuhn's to tight comparative statics: when orthodoxy is visible, contrarian motives move thresholds in predictable ways and—in the baseline-majority region—make additional information gathering privately worthwhile.

8 Discussion and Concluding Remarks

Our analysis shows how embedding contrarian motives into a sequential learning environment reshapes both information acquisition and welfare. On the positive side, nonconformist preferences expand the set of public beliefs under which agents find it worthwhile to experiment, thereby sustaining informative actions and counteracting premature herding. At the same time, these preferences lower chosen precision conditional on investing, reflecting the substitution of "being different" for "being accurate." The welfare consequences follow an inverted-U pattern: mild contrarianism

[14] Operational proxies for p_t include visible like/share counts, trending labels, or sorted feeds. Our comparative statics speak to these UI "knobs" even in non-algorithmic feeds because they work through observability, not recommendation logic.

is socially beneficial, but excessive contrarianism drives agents to act against already-informative public beliefs, generating correctness losses and inefficient experimentation.

Beyond the formal results, the framework sheds light on broader phenomena. In scientific communities, priority races and reputational incentives make standing apart from orthodoxy valuable, but contrarianism taken too far undermines shared standards of evidence. The same logic applies in cultural diffusion, where nonconformist motives keep diversity alive yet risk fragmentation. Our model thus offers a disciplined way to analyze when minority-seeking stabilizes experimentation and when it turns into destructive contrarianism.

The simplicity of our specification suggests several extensions. Allowing heterogeneity in contrarian intensity k—possibly private information—would capture richer identity-based motives and may generate assortative dynamics or segmentation. Introducing networks of observation, or continuous-time arrivals, could further align the model with empirical settings such as scientific collaboration or social media. Finally, linking contrarianism with forward-looking incentives (as in reputation models) may illuminate how agents balance distinctiveness with credibility.

Taken together, the results highlight a central message: contrarian motives, when moderate, improve learning by counteracting conformity, but when extreme, undermine the informativeness of actions. This trade-off connects the sociology of science with the economics of cascades, offering both a positive framework and a normative lens for thinking about conformity, diversity, and the dynamics of ideas.

A Additional proofs for Section 4

A.1 Proof of Proposition 4.3

The objective $G_t(\mu_t, \rho; p_t, k) - \frac{c}{2}\rho^2 - F\mathbf{1}\{\rho > 0\}$ is continuous on $[0, \infty)$ and goes to $-\infty$ as $\rho \to \infty$, so a maximizer exists. In the Gaussian case define

$$s_t^\star(\rho) = \tfrac{1}{2} + \frac{\tau_t - L_t}{\rho}, \qquad z_0(\rho) = s_t^\star(\rho)\sqrt{\rho}, \quad z_1(\rho) = (s_t^\star(\rho) - 1)\sqrt{\rho}.$$

Then

$$\Pr(a_t = 1 \mid \theta = 1) = 1 - \Phi\big(z_1(\rho)\big), \quad \Pr(a_t = 0 \mid \theta = 0) = \Phi\big(z_0(\rho)\big),$$

and $\mathbb{E}[1 - p_t(a_t)] = p_t + (1 - 2p_t)\Pr(a_t = 1)$. Differentiating gives

$$\frac{\partial}{\partial\rho} G_t(\mu_t, \rho; p_t, k) = \mu_t\,\varphi(z_1)(-z_1') + (1 - \mu_t)\,\varphi(z_0) z_0' + k(1 - 2p_t)\,\frac{\partial}{\partial\rho}\Pr(a_t = 1),$$

with $z_0' = \frac{s_t^\star}{2\sqrt{\rho}} + \sqrt{\rho}\, s_t^{\star\prime}$, $z_1' = \frac{s_t^\star - 1}{2\sqrt{\rho}} + \sqrt{\rho}\, s_t^{\star\prime}$, and $s_t^{\star\prime} = -(\tau_t - L_t)/\rho^2$. At any interior maximizer $\rho_t^\star > 0$ the Euler condition $\partial_\rho G_t = c\rho_t^\star$ holds.

A.2 Proof of Proposition 4.4

Let $P_1(\rho) = \Pr(a_t = 1)$. If $p_t > \frac{1}{2}$ and the no-signal action is 1, then $P_1(0) = 1$ and

$$\mathbb{E}[1 - p_t(a_t)] - \mathbb{E}[1 - p_t(a_t)]\big|_{\rho=0} = (1 - 2p_t)\big(P_1(\rho) - 1\big) \geq 0.$$

Thus the objective inside the $\sup_{\rho \geq 0}$ has increasing differences in (ρ, k), and the supremum $\Phi_t(\mu_t, p_t; k)$ is weakly increasing in k (the case $p_t < \frac{1}{2}$ is symmetric). Berge's maximum theorem yields continuity and the interval property of the investment region.

A.3 Proof of Proposition 5.3

For each $x = (\mu_t, p_t, k)$, the feasible set $[0, R]$ with R large ensures compactness; the objective is upper semicontinuous in ρ (downward jump only at 0) and continuous in (ρ, x) away from 0. The maximum theorem gives upper hemicontinuity of $\Gamma(x)$ and continuity of the value. Under uniqueness, strict SOC, and a nonsingular derivative of the FOC, the implicit function theorem yields local continuity/differentiability of $\rho_t^\star(x)$.

B Robustness: popularity–belief proximity

This lemma is *not used* in the proofs above; it only records that when p_t and μ_t are ε-close, comparative-statics directions coincide.

Lemma B.1. *Fix $k \geq 0$. If $|p_t - \mu_t| \leq \varepsilon$ and $\left|\mu_t - \frac{1}{2}\right| > \varepsilon$, then* $\operatorname{sign}\left(c_t - \frac{1}{2}\right) = \operatorname{sign}\left(\tilde{c}_t - \frac{1}{2}\right)$ *where $c_t = \frac{1}{2} + k(p_t - \frac{1}{2})$ and $\tilde{c}_t = \frac{1}{2} + k(\mu_t - \frac{1}{2})$.*

Proof. Immediate from $|c_t - \tilde{c}_t| \leq k\varepsilon$ and the triangle inequality around $\frac{1}{2}$. □

References

Acemoglu, D., Dahleh, M. A., Lobel, I., and Ozdaglar, A. (2011). Bayesian learning in social networks. *Review of Economic Studies*, 78(4):1201–1236. 2

Akerlof, G. A. and Kranton, R. E. (2000). Economics and identity. *Quarterly Journal of Economics*, 115(3):715–753. 3

Ali, S. N. and Kartik, N. (2012). Herding with collective preferences. *Economic Theory*, 51(3):601–626. 3

Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817. 2

Bénabou, R. and Tirole, J. (2011). Identity, morals, and taboos: Beliefs as assets. *Quarterly Journal of Economics*, 126(2):805–855. 3

Bernheim, B. D. (1994). A theory of conformity. *Journal of Political Economy*, 102(5):841–877. 3

Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026. 2

Bohren, J. A. (2016). Informational herding with model misspecification. *Journal of Economic Theory*, 163:222–247. 2

Bohren, J. A., Imas, A., and Rosenberg, M. (2019). The dynamics of discrimination: Theory and evidence. *American Economic Review*, 109(10):3395–3436. 2

Chamley, C. and Gale, D. (1994). Information revelation and strategic delay in a model of investment. *Econometrica*, 62(5):1065–1085. 3

Dasaratha, K. and He, K. (2020). Network structure and naïve sequential learning. *Theoretical Economics*, 15(2):643–689. 2

Dasgupta, A., Prat, A., and Verardo, M. (2011). The price impact of institutional herding. *The Review of Financial Studies*, 24(3):892–925. 3

Eyster, E., Rabin, M., and Weizsäcker, G. (2014). An experiment on social learning. *Journal of the European Economic Association*, 12(4):1143–1172. 2

Frick, M., Iijima, R., and Ishii, Y. (2020). Misinterpreting others and the fragility of social learning. *Econometrica*, 88(6):2281–2328. 2

Golman, R., Hagmann, D., and Loewenstein, G. (2017). Information avoidance. *Journal of Economic Literature*, 55(1):96–135. 3

Kartik, N. and Van Weelden, R. (2020). Informational herding with model uncertainty. *American Economic Review*, 110(12):3859–3893. 2

Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press, Chicago, 2nd edition. 13

Lu, F., Meyer, F., and Rosenbaum, F. (2019). Belief-dependent information acquisition in markets. Working paper; update authors, venue, and link if you have them. 3

Planck, M. (1949). *Scientific Autobiography and Other Papers*. Philosophical Library, New York. 13

Smith, L. and Sørensen, P. (2000). Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398. 2

Smith, L. and Sørensen, P. N. (2011). Observational learning. In Durlauf, S. N. and Blume, L. E., editors, *The New Palgrave Dictionary of Economics*. Palgrave Macmillan, London. Online edition. 2

Touboul, J. D. (2019). The hipster effect: When anti-conformists all look the same. *Discrete and Continuous Dynamical Systems–Series B*, 24(8):4379–4415. 3

Veldkamp, L. L. (2011). *Information Choice in Macroeconomics and Finance*. Princeton University Press, Princeton, NJ. 3

Vives, X. (2008). *Information and Learning in Markets: The Impact of Market Microstructure*. Princeton University Press, Princeton, NJ. 3