

Informational Herding, Optimal Experimentation, and Contrarianism

LONES SMITH

University of Wisconsin

PETER NORMAN SØRENSEN

University of Copenhagen

and

JIANRONG TIAN

University of Hong Kong

First version received August 2005; Editorial decision August 2020; Accepted January 2021 (Eds.)

In the standard herding model, privately informed individuals sequentially see prior actions and then act. An identical action herd eventually starts and public beliefs tend to “cascade sets” where social learning stops. What behaviour is socially efficient when actions ignore informational externalities? We characterize the outcome that maximizes the discounted sum of utilities. Our four key findings are: (1) cascade sets shrink but do not vanish, and herding should occur but less readily as greater weight is attached to posterity. (2) An optimal mechanism rewards individuals mimicked by their successor. (3) Cascades cannot start after period one under a signal log-concavity condition. (4) Given this condition, efficient behaviour is contrarian, leaning against the myopically more popular actions in every period. We make two technical contributions: as value functions with learning are not smooth, we use monotone comparative statics under uncertainty to deduce optimal dynamic behaviour. We also adapt dynamic pivot mechanisms to Bayesian learning.

Key words: Herding, Mimicking, Contrarian, Cascade, Efficiency, Monotonicity, Logconcavity

JEL Codes: C73, D82, D83.

1. INTRODUCTION

“What we have here is failure to communicate.”—Cool Hand Luke (1967)

In the benchmark ideal, markets fully aggregate dispersed information. When this does not happen, economics can be subtle and surprising. For instance, the possible collapse of trade with asymmetric information in Akerlof (1970) made sense of warranties and paved the way to the no-trade theorems in finance. In the last quarter century, the most cited new example of the failure of information communication has been informational herding. Here, the problem

The editor in charge of this paper was Dimitri Vayanos.

arises when privately informed and like-minded individuals sequentially act, after seeing prior actions. When action menus are *finite*—like investing or not investing—different private signals are pooled into the same action, and as a result, the learning process eventually breaks down in a striking way: the public history of past actions overwhelms all later private signals and a *herd* on some action starts.¹ In the herding model, everyone Bayes-updates the *public belief*—the posterior given the public history—with their endowed private signal. If private signals are boundedly strong, then strong enough public beliefs overwhelm them. Bikhchandani, Hirshleifer, and Welch (BHW) called this a *cascade*, and it guarantees a herd. In a cascade, actions reflect no new private signals, and so the public belief is no longer updated, in a failure of information aggregation.

This article optimally resolves the failure to communicate in informational herding: we formulate and solve the welfare problem for this paradigm. To do so, we set up the experimentation exercise by an infinite-lived planner who devises individual choice rules, *i.e.*, how to map private signals to observable actions. We also introduce a pivotal transfer scheme that decentralizes our solution. This is the first solution of the social planner's problem for the herding problem, amidst a vast literature on this topic. Our solution radically generalizes Banerjee's original 1992 conjectured solution.

This article then derives long-run and short-run characterization results of the planner's solution. We first show that the cascade set of public beliefs shrinks as the planner places higher weight on the welfare of future individuals. For our second major finding, we assert that the social planner should reduce the mimicry chance at the margin: Specifically, we show that individuals should act in this *contrarian* way: the more likely is a state, the more individuals should lean against the actions optimal in that state. This is optimal whenever the private signals obey an intuitive and often met condition that we describe below in posterior monotonicity (PM). This result offers clear behavioural predictions that apply in every period.

Our article makes four distinct novel contributions to dynamic information economics.

1. **OPTIMAL DYNAMIC INFORMATION DESIGN.** Proposition 1 derives a new index rule for the social planner's dynamic experimentation model. Individuals should take the action with the highest social value—namely, the weighted average of the individual's private value of the action and the informational gain to society. Since social values remain linear in private beliefs, the planner's payoff frontier is also piecewise linear in this private type. So, the planner recommends each action for an interval of private beliefs. But, his desire to signal information entails current sacrifices. An agent might optimally take myopically dominated actions, or take actions in a myopically suboptimal order (Appendix G.1), such as the low action for high beliefs and the high action for low beliefs!

We then practically decentralize this behavioural rule with a pivotal transfer scheme that affects the slope and intercept of the agent's private value function. Proposition 3 devises a Vickrey–Clarke–Groves mechanism that punishes mimickers and rewards anyone mimicked. This one-stage look-ahead scheme can internalize the informational externality because the successor's action is informative of the state. While this socially optimal incentive scheme is new, it is reminiscent of how academia rewards authors for citations.

2. **CASCADE SETS SHRINK BUT DO NOT VANISH WITH PATIENCE.** The planner's experimentation solution has a simple long-run takeout message: although cascade sets shrink, they do not vanish, so cascading remains socially efficient, only for a smaller set of public beliefs when posterity matters more (Proposition 2). To see why, note that learning shuts down in a cascade if and

1. Respectively, these are insights of Banerjee (1992) and Bikhchandani *et al.* (1992) (or BHW), and Smith and Sørensen (2000) (or SS). Chamley (2004) is an excellent distillation.

only if the planner's value function coincides with the myopic value function (Claim 8). But as the discount factor increases, optionality is more valuable, and the value function naturally rises. Hence, the static and planner's value functions coincide for fewer public beliefs, and the cascade set shrinks (see Figure 3). Finally, the planner's solution inherits from the original herding model that public beliefs converge to a cascade limit, and actions converge to a herd. The probability of a herd on an ex post suboptimal action falls to zero as the discount factor rises to one.

3. POSTERIOR MONOTONICITY PRECLUDES CASCADES. We next prune problematic updating behaviour from the herding model. Assume that two Bayesians Ike and Joe share a prior belief on the high state. Ike acts on the basis of his private belief, and Joe then sees this action choice, aware of Ike's Bayes-rationality. *Posterior monotonicity* asserts:

Prior belief rises \rightarrow Joe's posterior belief rises, conditional on a given action by Ike. (PM)

While (PM) is a compelling property, it can fail in a herding model since actions generate endogenous signals: for at higher prior beliefs, Ike takes any action for less favourable private signals, and so his action less strongly endorses the high state. For some signal distributions, this swamps the direct effect of a higher prior public belief.

Proposition 4 establishes that a log-concavity condition (LC) met by standard continuous signal distributions is equivalent to (PM). (PM) failures play a pivotal role in many applications of the herding model, as SS showed that *delayed cascades* (i.e. starting after period one) can arise only when (PM) fails. In other words, the discrete signal distribution in BHW was not merely illustrative but forced the delayed cascades. Condition (LC) rules out delayed cascades.

4. CONTRARIANISM. We conclude by exploring how the planner skews action choices at the margin. In the standard selfish herding model, everyone is indifferent between adjacent actions at fixed cutoff posterior beliefs. Proposition 5 deduces from (PM) that contrarian behaviour is efficient—the planner's cutoff private belief rises in the public belief. This finding arises in every period, and so is testable, such as done in Çelen and Kariv (2004). So a high action should be discouraged more as the public belief rises. Intuitively, at a higher public belief, the lower action is less likely, and so is more of a surprise. Thus, it is informationally more valuable. But when (PM) fails, the logic fails. At the higher public belief, with fixed cutoff, the posterior actually falls, which is already the desired effect—the planner needs not add to that.

RELATED LITERATURE. Banerjee (1992) first proposed a remedy for the social learning externality: conceal early actions. Our planner could choose to imitate concealment by dictating the resulting signal-to-action maps, but ignoring available information is in general suboptimal from the planner's own single-person decision problem perspective.² Centrally planned social learning is a topical and important problem, and new optimal mechanisms have recently been explored in applied settings.³

(a) ACTIONS AS SIGNALS. Our informational design problem may resemble the problem at the core of Bayesian persuasion (Kamenica and Gentzkow, 2011): in each period, our planner precommits to a state-dependent distribution over the actions that carry public signals to successive agents. In case 2 of the example in Section 2, the Professor partitions his private signals into one of two messages—one for each action—to maximize his Student's value. So, a Bayesian persuasion problem transpires in every period of our model, but where (1) the objective function is the

2. Doyle (2010) explores a planner's problem in Chamley and Gale's endogenous-timing herding model.

3. Actions include an unobserved information acquisition decision in Glazer *et al.* (2015).

endogenous value function and (2) the distribution over public signal realizations is constrained by the exogenously given private signal distribution and the number of available private actions.

[Dow \(1991\)](#) first attacked this type of problem, in which an observed action summarizes a private signal: a consumer observes a price realization, but in the next period can only recall its partition interval. In the second and final period, another price realization is seen, and a choice is made. The optimal determination of the first-period coarse price partition is like our planner's partition of signals. Like [Dow](#), our planner trades off the present and future, but the horizon is infinite, and he also struggles with an unknown state of the world.

(b) SOCIALLY EFFICIENT HERDING. Our planner's optimum is also a team equilibrium, where everyone maximizes the present value of utilities. As with our contrarianism, teams shy away more from more popular actions in [Vives's \(1997\)](#) continuum-agent Gaussian setting.⁴ [March and Ziegelmeyer \(2020\)](#) experimentally find evidence of contrarian behaviour among altruistic agents.⁵ Our contrarianism relates to the excessive imitation found in [Eyster and Rabin \(2014\)](#), but now in the traditional finite action herding setting. It explains why agents in herding models may efficiently excessively rely on their own signals, without assuming irrationality ([Eyster and Rabin, 2010](#)).

(c) BAYESIAN EXPERIMENTATION. The bad herding outcome parallels the familiar failure of complete learning in the two-armed bandit ([Rothschild, 1974](#)). Yet, the analogy is puzzling: [Easley and Kiefer \(1988\)](#) prove that complete learning generically arises in experimentation problems with finite state and action spaces, but the herding outcome likewise arises in a model with finite actions and states. This puzzle is resolved by our experimenting planner choosing a continuously defined signal-to-action map.⁶

The article opens with a two-period example of our two substantive findings of the planner's model: shrinking cascade sets and contrarianism. It then explores in sequence the social planner's problem as an experimentation model; the implementation results from our welfare indexes; the subtleties of dynamic information revelation; the way that cascade sets shrink in patience; the novel monotone restriction on signals that precludes cascades, and then the contrarianism finding that exploits this. Many proofs are appendicized.

2. AN ILLUSTRATIVE TWO PERIOD EXAMPLE

Our article proceeds indirectly, using recursion and dynamic programming. But, we first explicitly solve a simple two-period example that captures both of our main predictions of the planner's problem: shrinking cascade sets and contrarianism. Assume that economic theory research fashion is captured by one of two unobserved states, either low-brow theory L or high-brow theory H . A Professor and a Student share a *prior belief* π on state H . Respectively, they observe conditionally independent draws σ_P, σ_S of a private signal, with cdf's $F^H(\sigma) = \sigma^2$ and $F^L(\sigma) = \sigma$, and densities $f^H(\sigma) = 2\sigma$ and $f^L(\sigma) = 1$, as in [Figure 1](#). Since the signal likelihood ratio $f^H(\sigma)/f^L(\sigma) = 2\sigma$ in favour of state H increases, higher signals σ lead to higher posterior beliefs p in state H . Also, low $\sigma > 0$ are arbitrarily powerful indicators of state L , but all $\sigma < 1$ have bounded power for H .

4. [Vives](#) always employs the normal learning model, ruling out results like ours on the distributional shape's importance. On the other hand, that model characterizes the long-run properties of learning by the speed with which the public precision approaches infinity. Our analysis offers no analogy.

5. In a related setting, [Medrano and Vives \(2001\)](#) describe behaviour that reveals less private information as "contrarianism." We find it more natural that contrarian behaviour leans against the public belief.

6. [Pastorino and Kehoe \(2011\)](#) seek monotonicity of the optimal rule in a dynamic setting similar to ours. In their paper, the experimenter is constrained to choose from a finite set of interval partitions.




FIGURE 1

The Selfish Professor (Section 2)

Notes: At left are the private signal densities on $[0, 1]$: The likelihood odds favouring state H rise from 0 to 2 as the signal σ rises. At right are \vee -shaped myopic (expected) payoffs $U(\pi)$ and value $V(\pi)$ for the selfish Professor in the public belief π . These apply to the selfish Student if plotted against continuation belief p . In the cascade set $[0, 1/3]$, signals are useless, and the value and expected payoff coincide, $V(\pi) = U(\pi)$. If $\pi > 1/3$, information is valuable, and so $V(\pi) = \frac{1}{2}(5\pi - 4 + 1/\pi) > U(\pi)$.

After seeing his signal, the Professor either starts a low-brow paper l or high-brow paper h . His Student then learns from his paper choice, and makes his own paper selection. Research pays 1 if the paper and state match, and -1 otherwise (Figure 1). If the Professor updates to the posterior belief q in state H , his *expected payoff* is $U(q) \equiv \max(2q - 1, 1 - 2q)$.

We compare two extreme motivations for the Professor: he selfishly only cares about his own expected payoffs, or he is entirely motivated by his Student's expected payoffs.

Case 1: The Selfish Professor. Assume the Professor writes paper h when state H is most likely—*i.e.*, for posterior beliefs $q \geq 1/2$. By Bayes rule, this happens when his posterior likelihood ratio of states H to L exceeds one, or $[f^H(\sigma)/f^L(\sigma)][\pi/(1-\pi)] \geq 1$. This happens for high private signals σ above a *selfish threshold signal* $\bar{\sigma}(\pi) \equiv (1-\pi)/(2\pi)$. For any prior belief $\pi < 1/3$ in state H , the threshold signal impossibly exceeds one—in this case, the Professor always writes paper l . This event when the prior belief overwhelms all private signals is called a *cascade*. Here, the Professor's (prior expected) *value*—or highest expected payoff—is $V(\pi) = U(\pi)$ when $\pi < 1/3$. Otherwise, his payoff in state H is ± 1 for signals $\sigma \gtrless \bar{\sigma}(\pi)$, and oppositely so in state L . All told, the value is therefore:

$$V(\pi) = \pi[1 - 2F^H(\bar{\sigma}(\pi))] + (1-\pi)[2F^L(\bar{\sigma}(\pi)) - 1] = \frac{1}{2}(5\pi - 4 + 1/\pi).$$

This value $V(\pi)$ is strictly convex on $(1/3, 1)$, as depicted in Figure 1. Since the Professor profits from his information here, we have $V(\pi) > U(\pi)$ on $(1/3, 1)$. Indeed,

$$V(\pi) - (1 - 2\pi) = (3\pi - 1)^2/(2\pi) > 0 \quad \text{and} \quad V(\pi) - (2\pi - 1) = (\pi - 1)^2/(2\pi) > 0.$$

Consider next the Student's choice. After seeing the Professor's paper genre l or h , the Student, respectively updates to the *continuation prior belief* $p = p_l(\pi)$ or $p_h(\pi)$. Provided the Professor sometimes writes either genre, we have $p_l(\pi) < \pi < p_h(\pi)$. The selfish Student likewise writes the low-brow paper l iff her signal $\tau < \bar{\sigma}(p) = (1-p)/(2p)$. Similarly, the Student secures an expected value $V(p)$ from the continuation belief p . Suppose that the Professor is in the cascade set, with

$\pi \leq 1/3$. Since $p = p_l(\pi) \leq \pi$, whenever the Professor writes a low-brow paper, the Student copies him.

Case 2: The Altruistic Professor. Consider next the opposite extreme when the Professor chooses his paper genre to maximize his Student's expected value. Since $\bar{\sigma}(1/3) = 1$, with a prior belief at or just below $1/3$, the selfish Professor always chooses the low-brow paper, which sends the student a useless signal. To help the Student, by informing him of high signals, the altruistic Professor therefore leans against the prevailing prior belief, by choosing a lower *altruistic threshold signal* $\hat{\sigma} < \bar{\sigma}(\pi)$. In other words, he writes the high-brow paper more often, yielding respectively lower continuation beliefs: $\hat{p}_l(\pi, \hat{\sigma}) < p_l(\pi | \bar{\sigma}(\pi))$ and $\hat{p}_h(\pi, \hat{\sigma}) < p_h(\pi | \bar{\sigma}(\pi))$. We can explicitly compute them:

$$\hat{p}_l(\pi, \hat{\sigma}) = \frac{\pi \hat{\sigma}^2}{\pi \hat{\sigma}^2 + (1 - \pi) \hat{\sigma}} < \pi < \hat{p}_h(\pi, \hat{\sigma}) = \frac{\pi [1 - \hat{\sigma}^2]}{\pi [1 - \hat{\sigma}^2] + (1 - \pi) [1 - \hat{\sigma}]}, \quad (2.1)$$

and the action chances are the denominators of (2.1). The Professor chooses $\hat{\sigma}$ to maximize the Student's expected value. As in the right panel of Figure 1, if $\hat{p}_l(\pi) \geq 1/3$, this value is:

$$E[V(P)|\pi, \hat{\sigma}] = E[(5P - 4 + 1/P)/2 | \pi, \hat{\sigma}] = \frac{5}{2}\pi - 2 + \frac{1}{2}E[(1/P)|\pi, \hat{\sigma}], \quad (2.2)$$

where we have used the fact that $E[P|\pi, \hat{\sigma}] = \pi$, by the law of iterated expectations.

To evaluate the expectation (2.2), we can rule out the case where the continuation beliefs obey $\hat{p}_l(\pi) > 1/3$ for any prior belief $\pi > 1/3$.⁷ Rather, the Professor optimally endows the Student with the two continuation beliefs $\hat{p}_l(\pi) \leq 1/3 < \hat{p}_h(\pi)$. Since the Student's value is $V(p) = 1 - 2p$ on $[0, 1/3]$, and $V(p) = 5p - 4 + 1/p$ on $[1/3, 1]$, its expectation is $E[V(P)|\pi, \hat{\sigma}]$

$$= [\pi \hat{\sigma}^2 + (1 - \pi) \hat{\sigma}] (1 - 2\hat{p}_l(\pi)) + (\pi [1 - \hat{\sigma}^2] + (1 - \pi) [1 - \hat{\sigma}]) (5\hat{p}_h(\pi) - 4 + 1/\hat{p}_h(\pi)) / 2.$$

If we substitute the continuation beliefs (2.1), this expression reduces to

$$E[V(P)|\pi, \hat{\sigma}] = V(\pi) + \frac{(3\pi - 1)[1 - \pi \hat{\sigma}^2 - \pi \hat{\sigma} - \pi] \hat{\sigma}}{\pi(1 + \hat{\sigma})} + \frac{\pi(1 - \pi) \hat{\sigma}^2 (1 - \hat{\sigma})}{\pi(1 + \hat{\sigma})}. \quad (2.3)$$

Taking logs of the second two terms, the first-order condition in $\hat{\sigma}$ then simplifies to

$$4\pi^2 \hat{\sigma}^3 + 2\pi(5\pi - 1) \hat{\sigma}^2 + 4\pi(2\pi - 1) \hat{\sigma} + 3\pi^2 - 4\pi + 1 = 0. \quad (2.4)$$

Let us first see what this says about the cascade set. At the highest cascade prior belief $\bar{\pi}$, the Professor optimally always chooses the low-brow paper. In other words, the altruistic threshold signal is $\hat{\sigma} = 1$. This implies

$$0 = 4\bar{\pi}^2 + 2\bar{\pi}(5\bar{\pi} - 1) + 4\bar{\pi}(2\bar{\pi} - 1) + 3\bar{\pi}^2 - 4\bar{\pi} + 1 = (5\bar{\pi} - 1)^2.$$

For an alternative insight here, starting just above the prior belief $\bar{\pi} = 1/5$, if the Professor almost always chooses the low-brow paper, the high-brow paper signals σ_P very close to one, with

7. For if so, by substituting from (2.1), the last term in (2.2) reduces to:

$$E[(1/P)|\pi, \hat{\sigma}] = 1/\pi + \frac{(1 - \pi)^2}{\pi} \frac{1 - \hat{\sigma}}{1 + \hat{\sigma}}.$$

This expectation is falling in the threshold signal, maximized at $\hat{\sigma} = 0$. But for positive low signal thresholds $\hat{\sigma} > 0$, the low-brow genre l conveys a very discouraging message to the Student, resulting in a low continuation belief $p_l(\pi) < 1/3$. Given this contradiction, the lower continuation belief must lie in the cascade set: $p_l(\pi) < 1/3$. But unless the higher continuation belief $p_h(\pi)$ exceeds $1/3$, the student always finds himself in the cascade set, and information is worthless. In this case, he earns payoff $U(\pi)$.

likelihood ratio near 2. Hence, this endows the Student with the continuation belief just above $1/3$. This in turn leaves the Student just outside his cascade set. In summary, *altruism shrinks the Professor's set of cascade beliefs from $[0, 1/3]$ to $[0, 1/5]$.*

Next, one might ask how altruism impacts the Professor's actions at the margin, as he slowly grows more confident in state H . Consider the Professor's posterior odds $2y = 2\pi\hat{\sigma}/(1-\pi) \geq 1/2$ for state H , when he is at the knife-edge. Substituting this expression into (2.4), the optimal posterior odds $2y(\pi)$ on the domain $1/5 \leq \pi \leq 1$ obey

$$4(1-\pi)^2 y(\pi)^3 + 2(5\pi-1)(1-\pi)y(\pi)^2 + 4\pi(2\pi-1)y(\pi) - (3\pi-1)\pi = 0. \quad (2.5)$$

We can fortunately factor this cubic to get

$$[2(1-\pi)y(\pi) + 3\pi - 1][2(1-\pi)y(\pi)^2 + 2\pi y(\pi) - \pi] = 0.$$

With a cubic derivative with positive lead coefficient, the SOC requires the first or third solution of the FOC. But the first solution is negative. We thus need the positive solution of the quadratic:

$$y(\pi) = \frac{\sqrt{\pi^2 + 2\pi(1-\pi)} - \pi}{2(1-\pi)}. \quad (2.6)$$

This yields the optimal altruistic threshold signal $\hat{\sigma}(\pi) = (1-\pi)y(\pi)/\pi$ falling from 1 to 0 as π increase from 0.2 to 1. The corresponding threshold posterior belief increases. So, the Professor leans more against the high action the higher is π , as his indifference posterior odds $y(\pi)$ are higher. We call this property of an optimal solution *contrarianism*.

Recalling (2.3), the altruistic Professor's value function is—as plotted in Figure 2—the myopic payoff $U(\pi)$ on $[0, 1/5]$ inside the cascade, and on $[1/5, 1]$ is:

$$E[V(P)|\pi, \hat{\sigma}(\pi)] = V(\pi) + \frac{1-\pi}{\pi} \cdot \frac{\pi(3-2\pi) - \sqrt{1-(1-\pi)^2}}{\pi + \sqrt{1-(1-\pi)^2}}.$$

In this article, we argue that these two basic insights—shrinking cascade sets and contrarianism—are robust to any infinite horizon model of herding in which the social planner partially discounts the utility of future decision makers, provided signals obey (PM).

A key dynamic feature of this two period example owes to its non-stationarity: when the altruistic or selfish Professor is not in a cascade set, the Student lands in a cascade set if he sees paper l , since its continuation belief is $p_l(\pi) < 1/3$. By contrast, in our stationary infinite horizon model in Section 3, the cascade set is constant in all periods, and cannot be entered if (PM) obtains (as holds in this example's signal).

3. THE FORWARD-LOOKING HERDING MODEL

We start with the standard herding model of Smith and Sørensen (2000) (SS) and show that its planner's problem exactly corresponds to a single agent experimentation model. An infinite sequence of *agents* (decision-makers) share a common 50–50 prior belief, for simplicity, over two payoff relevant *states* $\omega \in \{L, H\}$. They act in the order $n = 1, 2, \dots$

Agents share a state-dependent utility function $u_\omega(a)$ over actions $a \in \{1, \dots, A\}$. Action 1 is uniquely best in state L , and action A in H . Payoffs obey increasing differences: $u_H(1) - u_L(1) < u_H(2) - u_L(2) < \dots < u_H(A) - u_L(A)$. Action a has *myopic payoff*, or expected payoff:

$$\bar{u}(a, r) = (1-r)u_L(a) + ru_H(a), \quad (3.1)$$

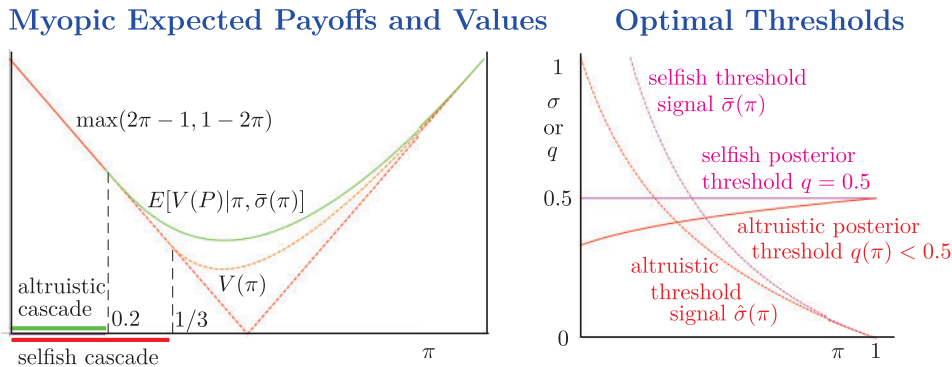


FIGURE 2

The Selfish and Altruistic Professors

Notes: The (solid) expected payoffs for the altruistic Professor exceed the value $V(\pi)$ of the selfish Professor, as he profits from the Student's signal too. A foretaste of our long run finding, the selfish cascade set is strictly smaller than the altruistic cascade set— $[0, 0.2] \subset [0, 1/3]$. At right, we depict our short run finding, contrarianism: as the prior chance of state H increases, the altruistic Professor writes paper l more often, and the threshold posterior q rises in π .

as a function of the probability r of state H . Since they may be useful for communication, we allow myopically dominated actions a , where $\bar{u}(a, r) < \sup_{\tilde{a}} \bar{u}(\tilde{a}, r)$ for all $r \in [0, 1]$.

The n th agent sees a random private signal σ_n about the state. We can identify this signal with the interim belief $\sigma_n = P(H|\sigma_n)$. The signals are i.i.d. across agents in each state $\omega = L, H$, with cdf F^ω . No signal perfectly reveals the state, so that F^H and F^L are mutually absolutely continuous with common support $\text{supp}(F)$. Because the signal is the interim belief, the derivative satisfies $dF^H/dF^L = \sigma/(1-\sigma)$. This is the “no-introspection property” deduced in our earlier SS paper—namely, that sampling an individual with signal σ should be just as informative as observing signal σ . Easily, this implies $F^H(\sigma) \leq F^L(\sigma)$, with strict inequality between the extremes of $\text{supp}(F)$. Signals are *unbounded* if their support $\text{supp}(F)$ contains 0 and 1, and *bounded*⁸ if $\text{supp}(F) \subseteq (0, 1)$. Also, some signals are informative: $\text{supp}(F)$ contains signals above and below $\sigma = 1/2$.

Individuals observe the history of actions but not private signals. Before choosing action a_n , the n th agent observes the history of $n-1$ predecessors' actions. Conjecturing their strategies, he can then deduce the updated public belief $\pi_n = P(H|a_1, \dots, a_{n-1})$ in state H . Combining his private signal σ with public belief π gives the posterior belief map:

$$r = R(\pi, \sigma) \equiv \frac{\pi \sigma}{\pi \sigma + (1-\pi)(1-\sigma)}. \quad (3.2)$$

This article explores welfare properties of this model: abstractly, the planner may modify how agents map private signals into actions, after any given history; in Section 6, we show how to implement this. A choice rule ξ prescribes some action $a = \xi(\sigma)$ for every signal σ .⁹ A strategy s_n for the n th agent assigns a choice rule to each action history of length $n-1$. The planner's preference depends on a discount factor $\delta \in [0, 1)$ that trades off payoffs earned by present and future agents. The planner chooses the strategy profile $s = (s_1, s_2, \dots) \in S$ to maximize the expected

8. In Section 2, $\text{supp}(F)$ contained 0 but not 1. For convenience, we henceforth omit this possibility.

9. If some signals σ have positive probability under F , optimization requires choice rules with a random action. For simplicity, we refer only to pure rules. Results remain valid for mixed choice rules.

present value of the utility stream $u_\omega(a_n)$, namely:

$$v_\delta(\pi) = \sup_{s \in \mathcal{S}} E[(1 - \delta) \sum_{n=1}^{\infty} \delta^{n-1} u_\omega(a_n)], \quad (3.3)$$

where we call v_δ the *value function*. The expectation is both over states ω and over the random process of private signals. The original herding model is the special case $\delta=0$.

4. DYNAMIC PROGRAMMING SOLUTION

We solve the social optimum (3.3) using dynamic programming. A stationary, or Markovian, *policy* assigns a choice rule ξ for every public belief π , our state variable. With rule ξ , action a happens with probability $\psi(a, \omega, \xi) = \int_{\xi^{-1}(a)} dF^\omega$ in state ω , and unconditionally with probability $\psi(a, \pi, \xi) = \pi \psi(a, H, \xi) + (1 - \pi) \psi(a, L, \xi)$ (slightly abusing notation). Action a results in *continuation public belief* $p(a, \pi, \xi) = \pi \psi(a, H, \xi) / \psi(a, \pi, \xi)$ when $\psi(a, \pi, \xi) > 0$. We call action a , and its continuation belief, *active* if $p(a, \pi, \xi) > 0$.

For any policy, starting at belief π , the continuation value of (3.3) is a function $v_\delta(\pi)$. By dynamic programming, the optimal (average present) value function v_δ solves the Bellman equation:

$$v(\pi) = \sup_{\xi \in \Xi} (T_\xi v)(\pi), \quad (4.1)$$

where the *policy operator* T_ξ maps any continuation value v into the current value, namely:

$$(T_\xi v)(\pi) = \sum_{a=1}^A \psi(a, \pi, \xi) [(1 - \delta) \bar{u}(a, p(a, \pi, \xi)) + \delta v(p(a, \pi, \xi))]. \quad (4.2)$$

Since the upper envelope of affine functions is convex, the value function v_δ solving (4.1) is convex, and therefore everywhere has a left and right derivative. Because the optimal strategies at beliefs 0 and 1, respectively entail taking actions 1 and A, Figure 3 arises:

Lemma 1 (Value Functions) *The value v_δ is bounded, continuous, and convex in public beliefs π , with extreme slopes $v'_\delta(0+) \geq u_H(1) - u_L(1)$ and $v'_\delta(1-) \leq u_H(A) - u_L(A)$.*

We do not solve the Bellman equation (4.1) as formulated, since it optimizes over policies. Rather, we work in value space. So inspired, recall the multi-armed bandit (Section 6.5 of Bertsekas, 1987), in which an experimenter each period chooses among finitely many actions, each providing a random and independent reward. Gittins (1979) solved for the optimal behaviour: Replace each action by a sure thing reward that subsumes its optionality; each period, one chooses the action with the highest such *Gittins index*.

We now argue that the planner's optimal policy admits an analogous index rule—even though the actions obviously do not have independent reward distributions: faced with the public belief π and private posterior r , the agent chooses the action a with the largest *welfare index* $w(a, \pi, r)$ —equal to the social payoff as privately gauged by the agent.

A convex function v has supporting *subtangents* $\tau(\pi, r)$ at public beliefs π , defined as functions of the posterior $r \in [0, 1]$, and uniquely defined when $v'(\pi)$ exists. We next design welfare indexes using these subtangents, and thereby implement the planner's outcome.

Proposition 1 (Optimal behaviour) For any public belief π , an agent with posterior belief r takes the action a with maximal *welfare index*

$$w(a, \pi, r) = (1 - \delta) \bar{u}(a, r) + \delta \tau(p(a, \pi, \xi), r), \quad (4.3)$$

for some supporting subtangent $\tau(p, r)$ to v at public belief p , evaluated at posterior r .

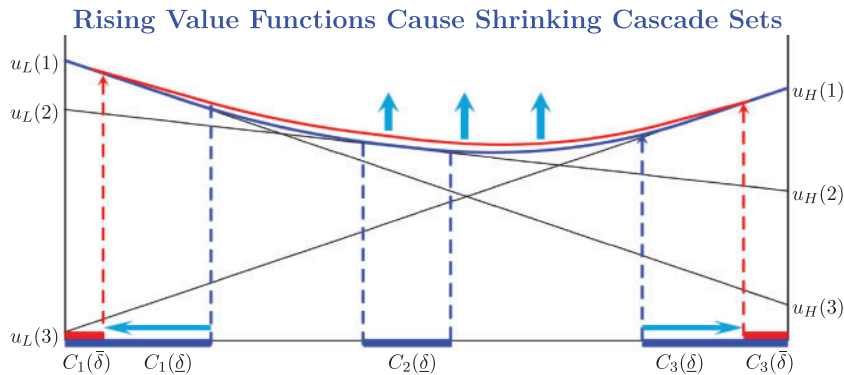


FIGURE 3

Myopic payoffs, the value function, and cascade sets for three actions

Notes: Cascades arise when the value function and myopic payoff coincide. Indeed, the value function strictly exceeds the myopic payoff when signals are valuable and can move beliefs, but coincide inside the cascade set intervals, because social learning stops. As the discount factor δ rises, the value increases, and the cascade sets therefore shrink. Interior cascade sets eventually vanish: $C_2(\bar{\delta}) = \emptyset$ and $C_2(\underline{\delta}) \neq \emptyset$ for $\underline{\delta} < \bar{\delta} < 1$, while extreme cascade sets shrink but never vanish: $C_1(\bar{\delta}) \supset C_1(\underline{\delta}) \neq \emptyset$ and $C_3(\bar{\delta}) \supset C_3(\underline{\delta}) \neq \emptyset$.

The value function subgradient $\tau(p, r)$ at a public belief p admits a useful economic interpretation: The marginal social benefit of a higher posterior is the slope $v'(p_a)$ at the continuation public belief $p_a = p(a, \pi, \xi)$. The privately informed agent thus assigns social value $\tau(p_a, r)$ to the continuation game after action a where followers act optimally at p_a . The welfare index (4.3) is a linear function of his posterior beliefs over the state of the world because both myopic payoffs and social incentives are.

Recall that the optimal strategy in the selfish herding model of SS was a simple *interval rule*: choose action a if one's posterior belief r lies in an interval I_a , where $\{I_a\}$ partitions $[0, 1]$. Actions with empty intervals are not taken. The rule may randomize at the *threshold* (boundary) θ_a between adjacent intervals I_a and I_{a+1} . Since the welfare index $w(a, \pi, r)$ in (4.3) is affine in r , interval rules remain socially optimal.¹⁰

Corollary 1 (Interval rules) An interval rule $\{I_a\}$ is optimal at any public belief π .

In the socially planned herding model, the communication value of actions can swamp myopic payoff considerations. Inspired by the search model of Dow (1991), we make two observations about the planner's optimal strategy reflecting this insight.¹¹

LESSON 1: MYOPICALLY DOMINATED ACTIONS MAY BE SOCIALLY VALUABLE.

Since more actions intuitively facilitates communication, myopically dominated actions might help. To see this, suppose that action A dominates $A - 1$ by ε in states L and H . We claim that for small $\varepsilon > 0$, action $A - 1$ is optimal with positive probability for some public beliefs and large $\delta < 1$. Intuitively, the static payoff losses can be made small, but the informational gain of an extra signal has boundedly positive value (proof in Section G.1).

LESSON 2: THE MYOPIC ACTION ORDER MIGHT NOT BE SOCIALLY OPTIMAL.

10. Dow (1991) derived an interval partition, albeit without our Bayesian binary state structure. Meanwhile, when $\delta = 1$, our Proposition 1 more roughly corresponds to the FOC of Dow's Proposition 2.

11. We generalize Dow's 1991 Proposition 2, which assumes perfect patience and a simple second-period value function. His Example 3 shows that a multiplicity of optimal solutions can arise in these problems.

The *natural order* entails using higher actions on higher intervals, if both are used. For high discount factors $\delta < 1$, the natural order need not be optimal (shown by example in Section G.1). But the action ordering is myopically optimal by our increasing differences assumption, and remains so with not too much concern for posterity. Define the gap $\Delta_a \equiv (u_H(a) - u_L(a)) - (u_H(a-1) - u_L(a-1))$ for actions $a \geq 2$. By increasing differences, $\Delta_a > 0$ for all $a \geq 2$. Next, define $\Delta = \Delta_2 + \dots + \Delta_A$ and $\underline{\Delta} = \min\{\Delta_2, \dots, \Delta_A\}$.

Corollary 2 For any discount factor $\delta < \underline{\Delta}/(\Delta + \underline{\Delta})$, the optimal rule obeys the natural action ordering. With two actions, this holds for $\delta < 1/2$.

5. CASCADE SETS: NON-EMPTY BUT SMALLER

For the selfish informational herding model, SS show that for bounded private signals, learning eventually ceases since the public belief π eventually lands in an absorbing state. The socially planned herding model has the same long run outcome. Let the *action cascade set* $C_a(\delta)$ be all public beliefs for which action a is optimal irrespective of the signal σ . In the Appendix, we piece together the dynamic learning story of how *public beliefs almost surely land in a cascade set*: since active learning ceases in $C_a(\delta)$, the value and myopic payoff coincide $v(\pi) = \bar{u}(a, \pi)$ in it; naturally, $C_a(\delta)$ is a closed interval, as in Figure 3. So *the value function is affine on cascade sets. Conversely, if the signal distribution has convex support and the natural action ordering is optimal, then the value function is affine only on cascade sets* (see Section B). So, the value is strictly convex iff active learning occurs.

For our first key finding of the article, we argue that the cascade phenomenon is efficient, but cascades happen too soon. For since the value function weakly increases in the discount factor (Claim 3), the value function coincides with the constant myopic payoff on weakly smaller cascade sets. In the perfect patience limit, the planner maximizes for the long run: these sets collapse to $\{0\}$, $\{1\}$, and incorrect herds no longer occur, as social learning is always correct (since a wrong point belief a.s. never happens).

Proposition 2 (Cascade Sets Shrink With Altruism, But Never All Vanish) (a) The value function strictly increases in δ for all public beliefs π outside cascade sets.
 (b) Cascade sets for an action are nested, strictly shrinking as $\delta < 1$ rises. For large $\delta < 1$, any $a \notin \{1, A\}$ has an empty cascade set, while $\lim_{\delta \rightarrow 1} C_1(\delta) = \{0\}$ and $\lim_{\delta \rightarrow 1} C_A(\delta) = \{1\}$.
 (c) A herd almost surely starts, and the probability it is incorrect vanishes as $\delta \uparrow 1$.

The shrinking of the cascade sets happens for a simple reason: since the planner is indifferent about experimentation at the edge of a cascade set, if he grows more patient, his value of information rises, and he strictly prefers to experiment; so the cascade set shrinks.

6. OPTIMAL DYNAMIC INFORMATION DESIGN

In a *team equilibrium* of our multistage game, everyone altruistically seeks to maximize the planner's payoff, taking other actions as given (Radner, 1962). We claim that *a social optimum is a team equilibrium for any discount factor $\delta < 1$* . To see why, suppose that all but one agent uses a sequentially rational optimal strategy s^* , but that someone has a strictly better reply $\hat{\xi}$ at a history in period n . Then the planner can improve his value at that history by *fully mimicking* this deviation, i.e., using rule $\hat{\xi}$ in the first period and then continuing with s^* as if s_n^* had been applied at stage n with this history (as the team would not have detected the deviation). This contradicts social optimality of strategy s^* .

While socially optimal behaviour is a team equilibrium, it need not constitute a Nash equilibrium. We wish to elicit altruistic behaviour from the selfish agents using a transfer scheme that only depends on the observed action history. By the index formula (4.3), we ultimately must award an agent the transfer $\delta \tau(\pi, r)/(1 - \delta)$. But, this depends on the unobserved private posterior belief r , and therefore is infeasible for the planner.

When the planner's policy prescribes an interval rule that does not swap the myopic interval order, it suffices to reward an agent just on the basis of his own action. For the planner can move the selfish agent's threshold between any two actions up (or down) by taxing (or subsidizing) the higher action. But transfers based on the agent's own action can never reverse the myopic action ordering, and so cannot implement the optimal action ordering. We solve this using a *pivot mechanism*—namely, one that rewards agents for their marginal contribution to social welfare—here, by changing the public belief.

In exchange for taking action a , an individual must be paid the *state ω contingent transfer* $t_\omega(a|\pi)$ equal to his successor's incremental value (seen in Figure 4):

$$t_\omega(a|\pi) = \frac{\delta}{1 - \delta} (v_\omega(p(a, \pi, \xi)) - v_\omega(\pi)), \quad (6.1)$$

defining the *state-contingent values* $v_L(p) \equiv \tau(p, 0)$ and $v_H(p) \equiv \tau(p, 1)$ at public belief p .

Since the transfers (6.1) depend on the unknown state ω , they are unavailable to the planner. Fortunately, since *future agents' choices reflect their information, and thereby the state*, we can surmount this hurdle. Let the next agent take the *least active action* \hat{a} for his lowest private signals $\sigma \leq \hat{\sigma}$. Because, as initially noted, $F^L(\hat{\sigma}) > F^H(\hat{\sigma})$ strictly inside $\text{supp}(F)$, action \hat{a} occurs more often in state L than state H . We can thereby emulate the incentive effect of the state-dependent transfer (6.1) by an *action transfer* $t(a, a')$ that depends just on the current and next actions, denoted a and a' :

Proposition 3 (Pivot Mechanism) The optimum (3.3) is implemented by a mechanism whose transfers depend on the public belief, and the actions of an agent and his successor.

Our proof in Section D simply asks whether the next agent takes the lowest active action \hat{a} or any other action $\neg \hat{a}$. The current agent earns the state ω contingent transfer $t_\omega(a|\pi)$ from taking

Value Function Tangents Determine Socially Optimal Transfers




FIGURE 4

Value function and optimal transfers

Notes: The subtangent $\tau(p, r)$ to the value function at public belief p yields the present value for someone with any posterior belief r (Proposition 1). Thus, higher posterior beliefs raise this value iff the value function slopes up. At extreme posteriors $r=0, 1$, the tangents at public beliefs $p=p_a$ or $p=\pi$ yield the state-contingent transfers $t_L(a|\pi)$ and $t_H(a|\pi)$ for action a in (6.1), directly proportional to (\propto) the respective indicated axis gaps. Proposition 3 implements these as functions of the successor's action (which reflects the unobserved state) and not the state.

action a if the following two linear equations in two unknowns have a solution:

$$t_{\omega}(a|\pi) = F^{\omega}(\hat{\sigma})t(a, \hat{a}) + [1 - F^{\omega}(\hat{\sigma})]t(a, \neg\hat{a}) \quad \text{for states } \omega = H, L. \quad (6.2)$$

This mechanism¹² implements the social best, as everyone earns his marginal contribution.

We next argue that with just two actions and the myopic action order, this pivot mechanism rewards individuals who are mimicked by successors. In other words, imitation is not only the best form of flattery, but also is socially optimal:

Corollary 3 (Mimicry is Optimally Rewarded) Assume the myopic action ordering with two actions. The pivot action transfers obey $t(a, a) > t(a, a')$ whenever $a' \neq a$ and neither public belief π nor the continuation belief $p(a, \pi, \xi)$ lie in the cascade set $C(\delta)$.

7. MONOTONE POSTERIOR BELIEFS

Before turning to our contrarianism result, we must dispense with a surprising Bayesian updating property—the failure of (PM). Our insight applies across information economics, as the desired new regularity property. Let Bayesians Ike and Joe share a common (prior) public belief π . Ike privately observes a conditionally independent private signal σ , and arrives at a private posterior belief $r = R(\pi, \sigma)$. Ike optimally chooses action $a \in \{1, 2, \dots, A\}$ when his belief r lies in the known interval I_a . Seeing his action, Joe infers that $r \in I_a$, and updates to a continuation public belief¹³ $p(a, \pi)$.

(PM) *Fixing any interval rule, any prior public beliefs $\pi' > \pi$, and any active action a (with $\psi(a, \pi), \psi(a, \pi') > 0$), continuation public beliefs are ranked $p(a, \pi') > p(a, \pi)$.*

So seeing the same action, a more optimistic prior leads to a more optimistic posterior belief. This may seem obvious, but it can fail. Joe's continuation public belief $p(a, \pi)$ averages over Ike's posteriors $R(\pi, \sigma) \in I_a$. As the prior π increases, Ike chooses action a for lower signals σ . In fact, for multinomial signals it can fail: suppose Ike takes the high action for realized signal σ_3 , but not for σ_1 or σ_2 . A slightly higher π might lead Ike to choose the high action also for σ_2 . This leads to a discontinuous drop in Joe's posterior.

Monotonicity is restored by an appropriate log-concavity assumption, (LC), applied to the log-likelihood ratio signal transformation $\ell = \Lambda(\sigma) \equiv \log(\sigma/(1-\sigma))$. Bayes rule (3.2) is then additive: $\Lambda(r) = \Lambda(\pi) + \Lambda(\sigma)$. Notice that $\Lambda(I_a)$ shifts $\Lambda(\pi)$ by a fixed interval of ℓ .

(LC) *The private signal distribution is atomless with convex support, and one state-contingent density for its log-likelihood ratio is strictly log-concave.*¹⁴

Property (LC) holds for both states, if it ever holds. For if the signal $\sigma = S(\ell) = e^{\ell}/(1+e^{\ell})$ has log likelihood ratio $\ell = \Lambda(\sigma)$, then by the chain rule, the state- ω -contingent density for the log-likelihood ratio is $\phi^{\omega}(\ell) \equiv f^{\omega}(S(\ell))S'(\ell)$. So ϕ^L and ϕ^H share a common support $\text{supp}(\phi)$.

12. Unlike the dynamic pivot mechanism in Bergemann and Välimäki (2010), each agent acts just once.

13. Let action a have chance $\psi(a, \pi)$ and lead to posterior public belief $p(a, \pi)$, fixing the agent's probabilistic map from posterior beliefs to actions (as π varies). In the interior of I_a , the agent adopts action a for sure, and at any boundary point r' of I_a the agent takes action a with fixed chance. Under (LC) stated below, boundary points occur with probability zero, and the exact mixing strategy becomes irrelevant.

14. Only atomless distributions with convex support can be log-concave, but some of the most familiar probability distributions have this property (see Marshall and Olkin (1979), Section 18.B.2.d).

Recalling Section 3, we have $(dF^H/dF^L)(\sigma) = \sigma/(1-\sigma)$ and thus $\phi^H(\ell)/\phi^L(\ell) = e^\ell$. Since $\log \phi^H = \log \phi^L + \ell$, we have ϕ^H log-concave iff ϕ^L is log-concave.

Proposition 4 (Posterior Monotonicity) Private signals obey (LC) iff (PM) holds.

Intuitively, once Bayes rule is additive, (LC) is equivalent to the monotone likelihood ratio property of private posterior $\Lambda(\pi) + \ell$ with respect to π , in turn equivalent to (PM). Log-linearity might then appear equivalent to posterior constancy, but log-linearity rules out unbounded beliefs, and the belief bound contributes to push up the posterior.¹⁵

The Professor–Student example in Section 2 offers an instructive knife-edge case for condition (LC). For its likelihood ratio $e^\ell = f^H(s)/f^L(s) = 2s$, with $f^L(s) = 1$, the above formula yields the log-linear state- L density $\phi^L(\ell) = \mathcal{S}'(\ell) = (1/2)e^\ell$. Consistent with Proposition 4, for any rule where threshold posterior $\pi 2\hat{\sigma}/[\pi 2\hat{\sigma} + (1-\pi)]$ is held constant, the continuation beliefs from (2.1) have $p_l(\pi, \hat{\sigma})$ constant and $p_h(\pi, \hat{\sigma})$ rising in π .

We noted above that (PM) fails for multinomial signals. We now give a *continuous density* example in which (PM) fails because (LC) fails. Choose $b \in (1/2, 1)$, and define the density $f(\sigma) = 1/(2-2b)$ on $[0, 1-b] \cup [b, 1]$, and $f(\sigma) = 0$ otherwise. Let $f^H(\sigma) = 2\sigma f(\sigma)$ and $f^L(\sigma) = 2(1-\sigma)f(\sigma)$. Suppose that action 2 is optimal for posterior beliefs over $1/2$. Given a public prior belief $\pi > b$, the posterior likelihood ratio after seeing action 2 is therefore:

$$LR(\pi) \equiv \frac{\pi}{1-\pi} \frac{\frac{1+b}{2} + \int_{1-\pi}^{1-b} \frac{\sigma}{1-b} d\sigma}{\frac{1-b}{2} + \int_{1-\pi}^{1-b} \frac{1-\sigma}{1-b} d\sigma}.$$

When $b > (1+2\sqrt{2})/7$, we verify that $LR(\pi)$ is decreasing on $(b, b+\epsilon)$ for some $\epsilon > 0$.¹⁶

A *delayed cascade* is one that starts after period one because public beliefs transition into a cascade set. In their “bounded beliefs example,” Smith and Sørensen (2000) found that this arises iff the map from posterior public belief after an action is not monotone in the prior belief. That is, landing in a cascade set requires that Condition (LC) fail:

Corollary 4 (No Delayed Cascades) Given (LC), there is no delayed cascade if $\delta = 0$.

Cascades arise in Bikhchandani *et al.* (1992) since they posit multinomial signals, violating (LC).¹⁷ A particularly popular application of the herding model uses binary signals, implying the failure of (PM).

8. CONTRARIANISM

We have seen that as the weight on posterity increases and δ increases, cascade sets shrink, and so strictly fewer public beliefs guarantee mimicry (Proposition 2). This captures a long run form of contrarianism: when people care more about posterity, they insist on more precise social

15. When π rises, the interval of private beliefs that maps to action a must shift down. Where this interval includes a belief bound, its downward shift either includes more signals at the top or fewer signals at the bottom. This effect creates the upward movement in the posterior upon observing action a .

16. Straightforward algebra provides $LR'(b) = (1+2b-7b^2)/(1-b)^4$, negative when $b > (1+2\sqrt{2})/7$.

17. Herrera and Hörner (2012) note for the binary action model that is equivalent to two immediate properties: increasing hazard ratio and increasing failure ratio. They copy arguments from Smith and Sørensen (2000) to prove this sufficient for ruling out cascades, but incorrectly claim that the property is necessary.

learning before blindly following the herd. We now formulate and prove a short run *contrarianism* that impacts behaviour in every period, and therefore is observably testable—such as using the methods developed in Çelen and Kariv (2004). We show that at the margin, as public beliefs shift weight towards the high state, it is socially efficient to lean more away from actions that are myopically better in that state.

Let $\theta_a(\pi)$ be the *optimal threshold* posterior between actions a and $a+1$ at the public belief π . To be precise, behaviour is *strictly contrarian* if, for any public beliefs $\pi' > \pi$ with the same optimal action ordering,¹⁸ any optimal thresholds $\theta(\pi) = (\theta_1(\pi), \dots, \theta_{A-1}(\pi))$ for posterior beliefs are strictly ranked: $\theta_a(\pi') > \theta_a(\pi)$ for all actions $a = 1, \dots, A-1$. In other words, fewer posteriors result in the higher action when the public belief is higher. Behaviour is *contrarian* if $\theta(\pi') \geq \theta(\pi)$.

Proposition 5 Let signals obey (LC). Behaviour is contrarian if $\delta \in [0, 1)$. If $\delta > 0$ and actions obey the natural order, then behaviour is strictly contrarian outside cascade sets.

While cascade sets shrink for all signal distributions (Proposition 2), contrarianism requires (PM). The example in Section 2 typifies this link, and in Appendix G.2, we see that contrarianism can fail without (PM).

For insight into the role of (PM), assume two active actions $a = 1, 2$ respectively optimal for posteriors in $[0, \theta]$ and $[\theta, 1]$, as in Figure 5. Seeing action a , we arrive at the continuation public belief $p(a, \pi, \theta)$. By Proposition 1, the optimal action a for any posterior r has the maximal welfare index $w(a, \pi, r)$. We use this to take an infinitesimal FOC. Given the Bellman operator in (4.2), slightly shifting posteriors from action 1 into action 2 by reducing the threshold by dr yields a net payoff gain (dr) times:

$$w(2, \pi, r) - w(1, \pi, r) = (1 - \delta)[\bar{u}(2, r) - \bar{u}(1, r)] + \delta[\tau(p(2, \pi, \theta), r) - \tau(p(1, \pi, \theta), r)]. \quad (8.1)$$

If the public belief π rises to π' , this adjusts difference (8.1) via the value function tangents.

Now, by (PM), the continuation beliefs rise: $p(a, \hat{\pi}, \theta) > p(a, \pi, \theta)$ for $a = 1, 2$. We can exploit a useful property of pairs of tangents to a convex function. Not only do the welfare indexes $w(1, p, r)$ and $w(2, p, r)$ coincide at the posterior belief threshold $r = \theta$, but more strongly, $w(1, p, r)$ crosses $w(2, p, r)$ from above at $r = \theta$. In other words, the net gain to taking the higher action grows in the posterior belief r . The posterior threshold $r = \theta$ where $w(1, \pi, \theta) = w(2, \pi, \theta)$ therefore rises, as desired: $\theta' > \theta$.

Lastly, strict contrarianism arises if $\delta > 0$ and the value function is strictly convex. For tangents in (4.3) have positive weight and strictly shift slope in π outside cascade sets.

9. CONCLUSION

This article makes major technical and four substantive contributions to social learning.

We formulate and solve the planner's problem for informational herding with a discounted concern for posterity (Sections 3–4). This has long remained open because it is technically challenging: we secure long run results by adapting optimal experimentation theory, and derive short run comparative statics without a derivative or a single-crossing property.

18. We have seen that the optimal action ordering generally depends on the public belief. But the threshold comparison $\theta(\pi) < \theta(\pi')$ is meaningful only when the same action ordering is optimal at π and π' , and that the set of active actions is identical. Abusing notion, call A the number of active actions. Fixing one interval rule, we re-label active actions so that higher actions are taken at higher signals.

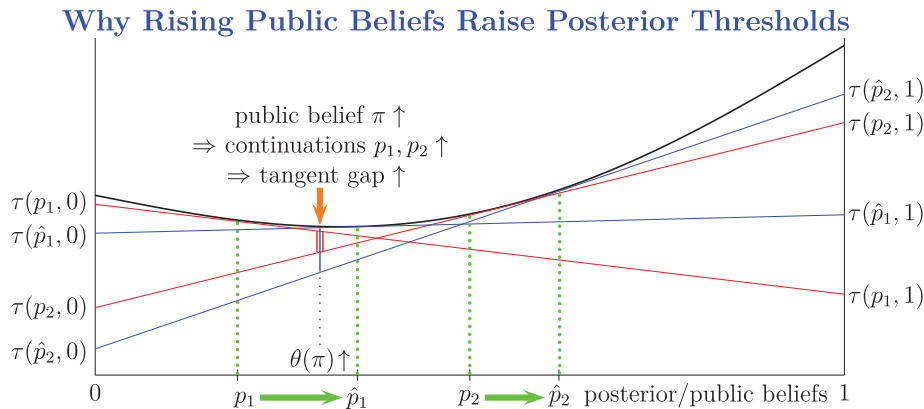


FIGURE 5

Contrarianism via comoving tangents

Notes: By (PM), when the prior public belief increases from π to $\hat{\pi}$, the continuation public beliefs $p_1 < p_2$ after actions $a=1$ and $a=2$ shift up to $\hat{p}_1 < \hat{p}_2$, respectively. Note that the gap between the (red) tangents to the value function at p_1 and p_2 is less than the analogous gap between the (blue) tangents at $\hat{p}_1 < \hat{p}_2$. The gap equals $(1-\delta)[u(2, \theta) - u(1, \theta)]/\delta$, by optimality equation (8.1). So to restore this equality, the threshold θ shifts higher for $\hat{\pi}$. All told, *contrarianism purely reflects how slopes of value function tangents co-move*. (See Section F.1.)

Section 5. The social optimality of herds has been questioned since Banerjee (1992): we prove that they remain socially optimal, but just that cascade sets shrink in the discount factor.

Section 6. We implement our planner's solution with a simple transfer scheme that rewards people who are followed. Since our agents need not introspect about others' preferences or optimal behaviour, this addresses a problem identified by Gagnon-Bartsch and Rabin (2016), that individuals might have limited understanding of others' optimal strategies.

Section 7. We derive a robust and frequently met new log-concavity property (LC) on signal distributions that notably precludes any delayed cascade in the standard herding model. Considering the central role of cascades in social learning, their fragility is important.

§8 We prove that, under property (LC), people should act in a contrarian way, leaning against increasingly popular actions *in every period*. Unlike the shrinking cascade set result, *this prediction is novel not merely for social learning, but also has no counterpart in the experimentation or Bayesian persuasion literatures*. It is a testable altruistic explanation for experimental herding work that agents excessively rely on their own signals.

In the popular symmetric two-action binary-signal model, the cascade set consists of two extreme public belief intervals, strictly shrinking with the discount factor. The efficient plan is simple: outside the cascade set, always follow the private signal. The constant rule renders it easy to compute the value function from the Bellman equation. This rule does not and cannot reveal any signs of contrarianism as we have defined it. But individuals in $C(0) \setminus C(\delta)$ should drop their private inclination to herd, and rather follow their signal.

Acknowledgments. We thank the editor and referees, as well as Abhijit Banerjee, Patrick Bolton, Katya Malinova, Meg Meyer, Piers Trepper, Chris Wallace, and seminar participants at MIT, Stockholm School of Economics, Stony Brook, Copenhagen, Lisbon, Bristol, Erasmus (Rotterdam), and Paris for comments on various versions. Smith gratefully acknowledges financial support for this work from National Science Foundation (NSF) grants SBR-9422988 and SBR-9711885, and Sørensen equally thanks the Danish Social Sciences Research Council. This article supersedes "Informational Herding and Optimal Experimentation" by the first two authors, based on Chapter 2 of Sørensen's 1996 thesis.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online. And the replication packages are available at <https://dx.doi.org/10.5281/zenodo.4443054>.

A. DYNAMIC PROGRAMMING PROOFS

In this section, we set up the dynamic programming optimization, and then use convex duality to prove that the solution in Proposition 1 and its properties are as claimed in Section 4.

VALUE FUNCTIONS: PROOF OF LEMMA 1. As Bayes rule does not identify $p(a, \pi, \xi)$ if $\psi(a, \pi, \xi) = 0$, we impose a weak refinement: $p(a, \pi, \xi) = R(\pi, \sigma)$ for some $\sigma \in \text{supp}(F)$.

We use the *Bellman operator* $T = \sup_{\xi \in \Xi} T_{\xi}$ from the RHS of (4.1). From (4.1) and (4.2), if $v \geq v'$ then $Tv \geq Tv'$. As is standard in discounted programs, T is a contraction, and so has a unique fixed point v_{δ} . This fixed point lies in the space of bounded, continuous, and convex functions. We show convexity. Since T is a contraction operator, it suffices that v convex implies Tv convex. Let $\pi_{\lambda} = \lambda\pi_1 + (1-\lambda)\pi_2$, where $\lambda \in (0, 1)$. Fix an optimal rule ξ mapping signals to actions at π_{λ} . Using Bayes rule, $p(a, \pi, \xi) = \pi \psi(a, H, \xi) / \psi(a, \pi, \xi)$, we get:

$$p(a, \pi_{\lambda}, \xi) = \frac{\lambda \psi(a, \pi_1, \xi)}{\psi(a, \pi_{\lambda}, \xi)} p(a, \pi_1, \xi) + \frac{(1-\lambda) \psi(a, \pi_2, \xi)}{\psi(a, \pi_{\lambda}, \xi)} p(a, \pi_2, \xi). \quad (\text{A.1})$$

The first (myopic) term in (4.2) at π_{λ} is the convex combination of the terms with π_1 and π_2 , as \bar{u} is linear in beliefs. As v is convex and (A.1) holds, the second (future) term obeys:

$$\psi(a, \pi_{\lambda}, \xi) v(p(a, \pi_{\lambda}, \xi)) \leq \lambda \psi(a, \pi_1, \xi) v(p(a, \pi_1, \xi)) + (1-\lambda) \psi(a, \pi_2, \xi) v(p(a, \pi_2, \xi)). \quad (\text{A.2})$$

Summing over actions $a = 1, \dots, A$ in (4.2) yields an upper bound on the Bellman operator:

$$Tv(\pi_{\lambda}) = T_{\xi} v(\pi_{\lambda}) \leq \lambda T_{\xi} v(\pi_1) + (1-\lambda) T_{\xi} v(\pi_2) \leq \lambda Tv(\pi_1) + (1-\lambda) Tv(\pi_2).$$

Let $U(\pi) = \max_a \bar{u}(a, \pi)$ denote the *myopic maximal expected payoff*. The bound on tangent slopes follows because actions 1 and A are respectively best in states L and H : $v(0) = u_L(1)$ and $v(1) = u_H(A)$, that the convex function v exceeds the myopic payoff U , and that $\bar{u}(1, r)$ and $\bar{u}(A, r)$ give the extreme slopes of U , by supermodularity.

Claim 1 (Iterates) $\{T^n U\}$ monotonely converges to the solution $v_{\delta} \geq U$ of (4.1).

Proof. Let rule $\tilde{\xi}$ a.s. choose the myopically optimal action, i.e., the maximizer over rules $\xi \in \Xi$ of $\sum_{a=1}^A \psi(a, \pi, \tilde{\xi}) [(1-\delta) \bar{u}(a, p(a, \pi, \tilde{\xi})) + \delta U(p(a, \pi, \tilde{\xi}))]$. Then $p(\tilde{\xi}(\sigma), \pi, \tilde{\xi}) = \pi$ a.s., and so the value $U(\pi)$. Optimizing over $\xi \in \Xi$, we get $TU(\pi) \geq U(\pi)$ for all π . By induction, $T^n U \geq T^{n-1} U$, a monotone sequence converging to a fixed point $v_{\delta} \geq U$. \parallel

Claim 2 (Slopes) $u_H(1) - u_L(1) \leq v'(0+) \leq v'(1-) \leq u_H(A) - u_L(A)$.

Proof. This follows graphically, as $v \geq U$ is convex, with $v(0) = U(0)$ and $v(1) = U(1)$. \parallel

Claim 3 (Weak Value Monotonicity) When $\delta_2 \geq \delta_1$, $v_{\delta_2}(\pi) \geq v_{\delta_1}(\pi)$ for all π .

Proof. Clearly, $\sum_{a=1}^A \psi(a, \pi, \xi) \bar{u}(a, p(a, \pi, \xi)) \leq \sum_{a=1}^A \psi(a, \pi, \xi) v(p(a, \pi, \xi))$ for any rule ξ and any function $v \geq \bar{u}$. If δ increases, then $T_{\xi} \bar{u}$ pointwise increases too, since more weight is placed on the larger component of the RHS of (4.2). By (4.1), $T\bar{u}$ is pointwise higher. Iterating this argument, $T^n \bar{u}$ is higher. Let $n \rightarrow \infty$, and appeal to Claim 1. \parallel

WELFARE INDEX CHARACTERIZATION: PROOF OF PROPOSITION 1. A convex function v is the upper envelope of its supporting tangent lines. Parameterized by their slope and intercepts, the *subtangent space* $T_v \subset \mathbb{R}^2$ is compact. Since \bar{u} and τ_a are affine functions, and since $p(a, \pi, \xi) = \int_{\xi^{-1}(a)} R(\pi, \sigma) dF^{\pi}$, we can exchange the order of summation and maximization to rewrite operator (4.2) as

$$(T_{\xi} v)(\pi) = \max_{(\tau_1, \dots, \tau_A) \in T_v} \sum_{a=1}^A \int_{\xi^{-1}(a)} [(1-\delta) \bar{u}(a, R(\pi, \sigma)) + \delta \tau_a(R(\pi, \sigma))] dF^{\pi}(\sigma). \quad (\text{A.3})$$

Exchange the sup in (3.3) with the max in (A.3) to get the planner's dual problem:¹⁹

$$v(\pi) = \max_{(\tau_1, \dots, \tau_A) \in \mathcal{T}_v^A} \sup_{\xi \in \Xi} \sum_{a=1}^A \int_{\xi^{-1}(a)} [(1-\delta)\bar{u}(a, R(\pi, \sigma)) + \delta\tau_a(R(\pi, \sigma))] dF^\pi(\sigma). \quad (\text{A.4})$$

The supremum over rules ξ in (A.4) entails allocating private signal σ to the action a where $(1-\delta)\bar{u}(a, R(\pi, \sigma)) + \delta\tau_a(R(\pi, \sigma))$ is maximal, and $r = R(\pi, \sigma)$. This yields the index expression (4.3). For a fixed rule ξ , (A.3) implies that τ_a is tangent to v at $p(a, \pi, \xi)$. \parallel

INTERVAL RULES: PROOF OF COROLLARY 1. The inner “sup” in optimization (A.4) asks, for every signal, which action is optimal. In this integral sum, for every signal, one takes the action with the highest index. This optimum is achieved by an interval rule. \parallel

PROOF OF LESSON 1: DOMINATED ACTIONS MAY BE SOCIALLY VALUABLE. A simple example suffices. Let $A-1$ be a dominated action with $u_L(A-1) > u_L(A) + \varepsilon^2$ and $u_H(A-1) = u_H(A) - \varepsilon$. Assume bounded signals: $\text{supp}(F) \subset (0, 1)$. Suppose first that $A-1$ were not available. By Claim 5(b), there is a cascade set $[\bar{\pi}, 1]$ for action A , where $\bar{\pi} < 1$. Since $v(\pi) = \bar{u}(A, \pi)$ for $\pi \in [\bar{\pi}, 1]$, and $v(\pi) > \bar{u}(A, \pi)$ for $\pi < \bar{\pi}$, the value function v is not locally linear at $\bar{\pi}$. Now, we make $A-1$ available, and check that it can improve this value. At belief $\bar{\pi}$, consider the rule that maps $\sigma \leq 1/2$ into action $A-1$, and $\sigma > 1/2$ into action A . This induces continuation public beliefs $p(A-1, \bar{\pi}, \xi) < \bar{\pi} < p(A, \bar{\pi}, \xi)$. Since the convex v is not locally linear at $\bar{\pi}$, the expected continuation value exceeds $v(\bar{\pi})$ by some $\eta > 0$. This policy change produces a myopic loss less than ε , beating the optimal policy when $\delta\eta > (1-\delta)\varepsilon$, i.e., for small enough ε . \parallel

NATURAL ACTION ORDER IF NOT TOO PATIENT: PROOF OF COROLLARY 2. By Proposition 1, it is optimal to choose the action with highest welfare index. Since $w(a, \pi, r)$ is linear in r , the natural order arises if $(\partial/\partial r)w(a, \pi, r)$ strictly rises in a . As v is convex, the slope of any subgradient line τ of v is sandwiched:

$$u_H(1) - u_L(1) \leq v'(0+) \leq \frac{\partial \tau}{\partial r} \leq v'(1-) \leq u_H(A) - u_L(A), \quad (\text{A.5})$$

recalling that the right derivative of a convex function is always uniquely defined. Hence:

$$\frac{\partial w(a+1, \pi, r)}{\partial r} - \frac{\partial w(a, \pi, r)}{\partial r} \geq (1-\delta)\Delta_{a+1} - \delta\Delta.$$

This is strictly positive when $\delta < \Delta_{a+1}/(\Delta + \Delta_{a+1})$. Finally, $\underline{\Delta} = \Delta_2$ for $A = 2$. \parallel

B. CASCADE SETS: PROOF OF PROPOSITION 2

The proof in this section characterizes stationary beliefs for the dynamic optimization.

Claim 4 (Strict Convexity of Values) *On cascade sets, the value function is affine and obeys $v(\pi) = U(\pi)$. Outside of cascade sets, $v(\pi) > U(\pi)$, and furthermore, if the signal support is convex and actions are naturally ordered, v is strictly convex.*

Proof. At any cascade public belief $\pi \in [\underline{z}, \bar{z}] \subset C_a(\delta)$, a.s. taking action a yields some state-contingent expected values v_L and v_H . So on $[\underline{z}, \bar{z}]$, $v(r) = (1-r)v_L + rv_H$ is affine.

Next, if π is not a cascade belief, by definition it is *not* optimal to induce one action a.s., whence $v_\delta(\pi) > U(\pi)$ if $\delta \in [0, 1)$. We next argue that it is strictly convex.

For a contradiction, let v be affine around a non-cascade belief $\hat{\pi} \notin C(\delta)$. Let $\hat{\xi}$ be an optimal rule mapping signals σ to actions a . Put $H(\pi) = \sum_a \psi(a, \pi, \hat{\xi}) w(a, \hat{\pi}, p(a, \pi, \hat{\xi}))$. First, $v(\hat{\pi}) = H(\hat{\pi})$, by (4.1) and Proposition 1. Second, H is affine because welfare indices are affine, and continuation public beliefs are a martingale (i.e. (A.1)). Third, $H(\pi) \leq v(\pi)$ from (A.4), as H employs both a particular rule $\hat{\xi}$ and particular tangents to v at $p(a, \hat{\pi}, \hat{\xi})$.

Since v is affine, $H(\pi) = v(\pi)$ around $\hat{\pi}$. By (A.4), $\hat{\xi}$ is optimal for π near $\hat{\pi}$ and $w(a, \hat{\pi}, r)$ is an optimal index function at π . Assume two active actions 1 and 2. As the actions are naturally ordered, $p(1, \hat{\pi}, \hat{\xi}) < p(2, \hat{\pi}, \hat{\xi})$ and $\bar{u}(2, r) - \bar{u}(1, r)$ is strictly increasing. So, the welfare indices $w(1, \hat{\pi}, r)$ and $w(2, \hat{\pi}, r)$ cross once. As the signal support is convex and both actions are active, the crossing θ uniquely fixes the optimal rule $\hat{\xi}$ at $\hat{\pi}$. But for public beliefs $\pi \neq \hat{\pi}$ near $\hat{\pi}$, the fixed private signal threshold rule $\hat{\xi}$ selects different actions for posteriors r near θ as π varies, since (3.2) is increasing in π . Contradiction. So the convex value function is not affine on any subinterval, and so is strictly convex. \parallel

19. As an aside, convex duality offers a computational strategy for solving the dynamic programming problem. In the iteration, given a value v_n , the next value v_{n+1} is obtained in principle by searching across all the possible rules. But convex duality suggests an alternative faster way to compute v_{n+1} : the required tangent space is merely the set of all the left and right derivative lines to v_n .

Claim 5 (Structure of Cascade Sets)

- (a) For discount factors $\delta \in [0, 1]$, we have $0 \in C_1(\delta)$ and $1 \in C_A(\delta)$, and $C(\delta) \neq [0, 1]$.
 (b) With bounded signals, $C_1(\delta) = [0, \underline{\pi}(\delta)]$ and $C_A(\delta) = [\bar{\pi}(\delta), 1]$ for $0 < \underline{\pi}(\delta) < \bar{\pi}(\delta) < 1$.
 (c) With unbounded signals, $C_1(\delta) = \{0\}$, $C_A(\delta) = \{1\}$, and $C_a(\delta) = \emptyset$ for $a \neq 1, A$.

Proof of (a). Action 1 is myopically strictly optimal if $\pi = 0$. Since it always updates to continuation belief $\pi = 0$, it is dynamically optimal for any discount factor $\delta \in [0, 1]$. Ditto for $\pi = 1$. As signals are valuable in the selfish problem, $\cup_{a=1}^A C_a(0) \neq [0, 1]$. \parallel

Proof of (b). Action 1 is strictly optimal at belief $\pi = 0$, and so selfishly optimal for $\pi \leq \pi'$, for some $\pi' > 0$. In particular, $\bar{u}(1, \pi) > \bar{u}(a, \pi) + \eta$ for all $a \neq 1$ for some $\eta > 0$, and for all $\pi \in [0, \pi'/2]$. No action can reveal a stronger private signal than any $\sigma \in \text{supp}(F) \subseteq [\sigma_0, \sigma_1] \subset (0, 1)$. So any initial belief π is updated to at most $\bar{p}(\pi) = \pi \sigma_1 / [\pi \sigma_1 + (1 - \pi)(1 - \sigma_1)]$. For π small enough, $\bar{p}(\pi) \in [0, \pi'/2]$ and $\bar{p}(\pi) - \pi$ is arbitrarily small. By continuity, $v_\delta(\bar{p}(\pi)) - v_\delta(\pi)$ is less than $\eta(1 - \delta)/\delta$ for small enough π . By the Bellman equation (4.1), any $a > 1$ is strictly suboptimal for such small beliefs. \parallel

Proof of (c). For unbounded signals, SS deduce $C_a(0) = \emptyset$ whenever $1 < a < A$, and $C_1(0) = \{0\}$ and $C_A(0) = \{1\}$. Cascade sets weakly shrink in δ by Claims 3 and 4. \parallel

Claim 6 (Cascade Sets Shrink at Interior Edges) For any discount factor $\delta > 0$ and action $a \in \{1, \dots, A\}$, if $\check{\pi} \in (0, 1)$ is an endpoint of the cascade set $C_a(0)$, then $\check{\pi} \notin C_a(\delta)$.

Proof. We focus on left endpoints, and so on an action $a > 1$. Let $\check{\pi} = \min C_a(0)$, with least posterior belief $\check{r} = R(\check{\pi}, \text{minsupp}(F)) < \check{\pi}$, for which (\dagger) : $\bar{u}(a - 1, \check{r}) = \bar{u}(a, \check{r})$. Put

$$w_{a-1}(r) = (1 - \delta)\bar{u}(a - 1, r) + \delta\tau(\check{r}, r) \quad \text{and} \quad w_a(r) = \bar{u}(a, r). \quad (\text{B.1})$$

As $\check{r} < \check{\pi}$, we have $\check{r} \notin C_a(0)$. Since $C_a(\delta) \subseteq C_a(0)$ by Claim 3, we have $\check{r} \notin C_a(\delta)$. By Claim 4, $\bar{u}(\check{r}, a) < v_\delta(\check{r}) = \tau(\check{r}, \check{r})$. So (B.1) and (\dagger) imply $w_a(\check{r}) < w_{a-1}(\check{r})$. If $\check{\pi} \in C_a(\delta)$, then $w_a(\check{r})$ is the welfare index at posterior \check{r} , contradicting Proposition 1. \parallel

Claim 7 (Continuations) Continuation beliefs lie in at most one cascade set if $\delta > 0$.

Proof. With unbounded signals, absent perfectly revealing signals, continuations never lie in a cascade set. Assume bounded signals. Let $\underline{\sigma} = \text{minsupp}(F)$ and $\bar{\sigma} = \text{maxsupp}(F)$. Assume two continuations $p_1 < p_2$ for some π lie in distinct cascade sets, $p_1 \in C_{a'}(\delta)$ and $p_2 \in C_{a''}(\delta)$, with $C_{a'}(\delta)$ below $C_{a''}(\delta)$ in $[0, 1]$. Then $p_1 \in C_{a'}(0)$ and $p_2 \in C_{a''}(0)$, by Proposition 2(b). Let $\pi' \equiv \max C_{a'}(0) \leq \min C_{a''}(0) \equiv \pi''$. Then $p_1 \leq \pi'$. Choose x_1, x_2 in $[\underline{\sigma}, \bar{\sigma}]$ with $R(\pi, x_1) = p_1$ and $R(\pi, x_2) = p_2$. As Bayes-updating commutes:

$$R(p_1, \bar{\sigma}) = R(R(\pi, x_1), \bar{\sigma}) = R(R(\pi, \bar{\sigma}), x_1) \geq R(R(\pi, x_2), x_1) \geq R(p_2, \underline{\sigma}) \geq R(\pi'', \underline{\sigma}) \geq R(\pi', \bar{\sigma})$$

and so $p_1 \geq \pi'$. Thus $p_1 = \pi' \equiv \max C_{a'}(0)$, which contradicts Claim 6. \parallel

Claim 8 (Strict Value Monotonicity) The value function strictly increases in $\delta \in [0, 1]$ outside the cascade sets: If $\delta_2 > \delta_1$, then $v_{\delta_2}(\pi) > v_{\delta_1}(\pi)$ for all public beliefs $\pi \notin C(\delta_2)$.

Proof. Let $\delta_2 > \delta_1$, and fix $\pi \notin C(\delta_2)$. If $\pi \in C(\delta_1)$, we're done, since $v_{\delta_1}(\pi) = U(\pi) < v_{\delta_2}(\pi)$. If $\pi \notin C(\delta_1)$, then the δ_1 -optimal rule ξ sometimes induces an action \hat{a} with continuation $p(\hat{a}, \pi, \xi) \notin C(\delta_1)$. [The public belief π is the average of continuations, each cascade set $C_a(\delta_1)$ is an interval, and at most one cascade set is hit, by Claim 7.] Then,

$$(1 - \delta_1)\bar{u}(a, p(a, \pi, \xi)) + \delta_1 v_{\delta_1}(p(a, \pi, \xi)) \leq (1 - \delta_2)\bar{u}(a, p(a, \pi, \xi)) + \delta_2 v_{\delta_2}(p(a, \pi, \xi))$$

for every action a by Claim 3, with strict inequality for some action \hat{a} , since $\delta_2 > \delta_1$ and $v_{\delta_2}(p(\hat{a}, \pi, \xi)) > v_{\delta_1}(p(\hat{a}, \pi, \xi))$. By (4.1) and (4.2), the δ_1 -optimal rule ξ strictly exceeds $v_{\delta_1}(\pi)$, for the discount factor δ_2 . Optimizing over δ_2 -rules, we have $v_{\delta_2}(\pi) > v_{\delta_1}(\pi)$. \parallel

Claim 9 (Shrinkage) With bounded beliefs, action cascade sets strictly shrink in δ .

Proof. Fix action \hat{a} and $0 < \delta_1 < \delta_2 < 1$. As $C_{\hat{a}}(\delta) = \{\pi | v_\delta(\pi) - \bar{u}(\hat{a}, \pi) = 0\}$ is closed by continuity, we prove that if the cascade set edge $\hat{\pi} \equiv \min C_{\hat{a}}(\delta_1)$, then $\hat{\pi} \notin C_{\hat{a}}(\delta_2)$ if $\hat{a} \neq 1$.

CASE 1: MULTIPLE OPTIMIZERS. Assume that at $\hat{\pi}$, some optimal rule for δ_1 uses actions other than \hat{a} with positive probability. Then some continuation beliefs fall outside $C(\delta_1)$ with positive probability. For if not, by Claim 7, all

continuations lie in the same cascade set; therefore, not playing \hat{a} incurs a myopic cost with no informational gain. As in the proof of Claim 6, $v_{\delta_2}(\hat{\pi}) > v_{\delta_1}(\hat{\pi}) = \bar{u}(\hat{a}, \hat{\pi})$, and therefore $\hat{\pi} \notin C(\delta_2)$.

CASE 2: A UNIQUE OPTIMIZER. Assume that the unique optimal rule for $\delta = \delta_1$ at belief $\hat{\pi}$ is almost surely to play \hat{a} . Choose a belief sequence $\pi_n \uparrow \hat{\pi} = \min C_{\hat{a}}(\delta_1)$, so that $\pi_n \notin C_{\hat{a}}(\delta_1)$. For each n , let $T_n = (\tau_a^n, a = 1, \dots, A)$ be optimal tangents in (A.4). For every n , $\pi_n \notin C_{\hat{a}}(\delta_1)$ implies the existence of action $\check{a}_n \neq \hat{a}$ and signal $\sigma_n \in \text{supp}(F)$:

$$(1 - \delta_1)\bar{u}(\check{a}_n, R(\pi_n, \sigma_n)) + \delta_1 \tau_{\check{a}_n}^n(R(\pi_n, \sigma_n)) \geq (1 - \delta_1)\bar{u}(\hat{a}, R(\pi_n, \sigma_n)) + \delta_1 \tau_{\hat{a}}^n(R(\pi_n, \sigma_n)).$$

Since $T_v^{[A]}$ and $\text{supp}(F)$ are compact and A is finite, there is a subsequence where T_n has limit $T^* = (\tau_a^*, a = 1, \dots, A)$, \check{a}_n has limit $\check{a} \neq \hat{a}$, and σ_n has limit $\hat{\sigma}$. Write $\hat{r} = R(\hat{\pi}, \hat{\sigma})$. By the Theorem of the Maximum, T^* is an optimal subgradient vector for $\hat{\pi}$, so that

$$(1 - \delta_1)\bar{u}(\check{a}, \hat{r}) + \delta_1 \tau_{\check{a}}^*(\hat{r}) \geq (1 - \delta_1)\bar{u}(\hat{a}, \hat{r}) + \delta_1 \tau_{\hat{a}}^*(\hat{r}). \quad (\text{B.2})$$

Since \hat{a} is uniquely optimal at $\hat{\pi}$, subgradient $\tau_{\hat{a}}^*$ must be identical to the myopic function $\bar{u}(\hat{a}, \cdot)$. By Claim 6, $\hat{\pi}$ is strictly inside $C_{\hat{a}}(0) \supset C_{\hat{a}}(\delta_1)$. So action \hat{a} is myopically strictly dominant, i.e., $\bar{u}(\hat{a}, \hat{r}) > \bar{u}(\check{a}, \hat{r})$. Then, (B.2) implies $\tau_{\check{a}}^*(\hat{r}) > \tau_{\hat{a}}^*(\hat{r})$. Hence, $(1 - \delta_2)\bar{u}(\check{a}, \hat{r}) + \delta_2 \tau_{\check{a}}^*(\hat{r}) > (1 - \delta_2)\bar{u}(\hat{a}, \hat{r}) + \delta_2 \tau_{\hat{a}}^*(\hat{r})$, as $\delta_2 > \delta_1$. Optimizing over tangents in (A.4) at δ_2 yields $v_{\delta_2}(\hat{\pi}) > \bar{u}(\hat{a}, \hat{\pi})$. So $\hat{\pi} \notin C_{\hat{a}}(\delta_2)$. \parallel

Claim 10 (The Perfect Patience Limit) Interior cascade sets are empty for large $\delta < 1$, while $\lim_{\delta \rightarrow 1} C_1(\delta) = \{0\}$ and $\lim_{\delta \rightarrow 1} C_A(\delta) = \{1\}$.

Proof. Fix a cascade set $C_a(\delta) \neq \emptyset$ for action $a \notin \{1, A\}$. It suffices to use a simple rule ξ taking action $a - 1$ for $\sigma \in I_{a-1} = [0, \theta]$ and otherwise action a , where both actions can happen: $0 < F(\theta) < 1$. By continuity of $p(a, \pi, \xi)$ in π on the compact subset $C_a(\delta) \subset (0, 1)$, for some $\varepsilon > 0$, after the good news $\sigma \in I_a$, public beliefs rise by $p(a, \pi, \xi) - \pi \geq \varepsilon$ for all $\pi \in C_a(\delta)$. Let $\pi'' = \max C_a(\delta)$ and define $[\pi', \pi''] \equiv [\pi'' - \varepsilon/2, \pi''] \cap C_a(\delta)$. Since the convex function $v_{\delta}(p) \geq \bar{u}(a, p)$ with equality iff $p \in C_a(\delta)$, by Claim 5(a), there exists $\eta > 0$ so that $\psi(a - 1, \pi, \xi)v_{\delta}(p(a - 1, \pi, \xi)) + \psi(a, \pi, \xi)v_{\delta}(p(a, \pi, \xi)) > v_{\delta}(\pi) + \eta$ on $[\pi', \pi'']$.

We claim that the interval $[\pi', \pi'']$ is excised from $C_a(\delta')$ for large enough $\delta' \in (\delta, 1)$. For if $v_{\delta'}(\pi') = \bar{u}(a, \pi')$, and thus in $[\pi', \pi'']$, switching from the cascade rule to ξ yields a continuation Bellman value in (4.1) at least η higher. For large $\delta' \in (\delta, 1)$, this gain exceeds any first-period loss, proving sub-optimality of the cascade rule at π' , and so in $[\pi', \pi'']$. After finitely many iterations, each slicing an $\varepsilon/2$ interval, $C_a(\delta)$ vanishes for large δ . By repeating this, for all $\varepsilon > 0$, $C_a(\delta) \cap [\varepsilon, 1 - \varepsilon]$ vanishes for large $\delta < 1$ near 1. \parallel

C. HERDING VIA BELIEF DYNAMICS: PROOFS

A herd obtains on action a at stage N if everyone henceforth choose action a . We argue in this section that a herd starts—for if not, beliefs could not converge, violating the martingale convergence theorem. Also, the limit is almost surely never fully wrong.

Claim 11 (Limit Beliefs) The public belief process $\langle \pi_n \rangle$ is a martingale unconditional on the state, converging a.s. to some limit r.v. π_{∞} . In state H , it a.s. belongs to $(0, 1]$.

Public beliefs tend to the cascade set, for if not, limit actions would be informative.

Theorem 6 (Limit Beliefs are Cascade Sets) The limit belief π_{∞} of the planner's problem has support in $C_1(\delta) \cup \dots \cup C_A(\delta)$, and so is concentrated on the truth for unbounded signals.

Proof. At a non-cascade belief π , at least two actions have positive probability. By Corollary 1, the highest such action is more likely in state H , and the least in state L . So, the continuation belief differs from π with positive probability. By the Markov-martingale process characterization in Appendix B of SS, this non-cascade belief $\pi \notin \text{supp}(\pi_{\infty})$. \parallel

Theorem 7 (Efficient Herds) A herd a.s. starts. For unbounded signals, it is on the ex post optimal action, and for bounded signals, it is incorrect with vanishing chance as $\delta \uparrow 1$.

We extend the “Overturning Principle” of SS to show that herds arise. Claim 13 below proves that actions $a' \neq a$ greatly move public beliefs π near the cascade set $C_a(\delta)$ —for such (unexpected) actions yield a first order myopic loss and second-order information gain. This is the contrapositive of: limit cascade (a.s. occurs by Theorem 6) \implies herd.

Proof of Theorem 7. Fix an optimal policy—a map Υ from public beliefs to rules, $\xi = \Upsilon(\pi)$. Let $\varepsilon > 0$ obey Claim 13. Define events $\mathcal{B}_n^a = \{\pi_n \text{ is } \varepsilon\text{-close to } C_a(\delta)\}$, $\mathcal{C}_n^a = \{\psi(a, \pi_n, \Upsilon(\pi_n)) < 1 - \varepsilon\}$, and $\mathcal{D}_{n+1} = \{|\pi_{n+1} - \pi_n| > \varepsilon\}$. Given $\mathcal{B}_n^a \cap \mathcal{C}_n^a$, Claim 13 (ii) yields $P(\mathcal{D}_{n+1} | \pi_n) \geq \varepsilon/A$. Then $\sum_{n=1}^{\infty} P(\mathcal{D}_{n+1} | \pi_1, \dots, \pi_n) = \infty$ given $\mathcal{B}_n^a \cap \mathcal{C}_n^a$ infinitely often (i.o.). By the Conditional Second Borel–Cantelli Lemma,²⁰ a.s. \mathcal{D}_n obtains i.o. given $\mathcal{B}_n^a \cap \mathcal{C}_n^a$ i.o. Since $\langle \pi_n \rangle$ a.s. converges by Claim 11, \mathcal{D}_n i.o. or $\mathcal{B}_n^a \cap \mathcal{C}_n^a$ i.o. have chance 0.

Let \mathcal{E}_a be the event that $\langle \pi_n \rangle$ has a limit in $C_a(\delta)$, intersected with the probability-one event that $\mathcal{B}_n^a \cap \mathcal{C}_n^a$ occurs finitely often. By definitions of \mathcal{B}_n^a and \mathcal{D}_{n+1} , convergence to $C_a(\delta)$ implies that eventually $\mathcal{E}_a \subset \mathcal{B}_n^a \setminus \mathcal{D}_{n+1}$. Since $\mathcal{B}_n^a \cap \mathcal{C}_n^a$ occurs finitely often, eventually $\mathcal{E}_a \subset \mathcal{B}_n^a \setminus (\mathcal{C}_n^a \cup \mathcal{D}_{n+1})$. By Claim 13(i), $\mathcal{B}_n^a \setminus \mathcal{C}_n^a$ implies that every $a' \neq a$ leads to \mathcal{D}_{n+1} , so actions $a' \neq a$ can occur only finitely often. So, action a is eventually taken on \mathcal{E}_a . By Claim 11 and Theorem 6, event $\bigcup_{a=1}^A \mathcal{E}_a$ has chance one. So, a herd starts. \parallel

Claim 12 (Bad Herds) *The incorrect herd chance vanishes as $\delta \uparrow 1$, for bounded signals.*

Proof. As $((1 - \pi_n)/\pi_n)$ is a convergent, non-negative martingale in state H , by Fatou’s Lemma, $E[\lim_{n \rightarrow \infty} (1 - \pi_n)/\pi_n] \leq (1 - \pi_0)/\pi_0$. Let the boundary of $C_1(\delta)$ have likelihood ratio $M < \infty$. Then, $\lim_{n \rightarrow \infty} (1 - \pi_n)/\pi_n > M$ with chance at most $1/M$. Since $M \uparrow \infty$ as $\delta \uparrow 1$ by Proposition 2(b), $\lim_{n \rightarrow \infty} \pi_n \in C_1(\delta)$ with a vanishing chance as $\delta \uparrow 1$. \parallel

We now argue that near cascade sets, some action boundedly moves the beliefs.

Claim 13 (A New Overturning Principle) *For $\delta \in [0, 1)$, let $C_a(\delta) \neq \emptyset$ for action a . Then there exists $\varepsilon > 0$ and an ε -neighbourhood $K \supset C_a(\delta)$, s.t. $\forall \pi \in K \cap (0, 1)$, either:*

- (i) $\psi(a, \pi, \Upsilon(\pi)) \geq 1 - \varepsilon$, and $|p(a', \pi, \Upsilon(\pi)) - \pi| > \varepsilon$ for all active actions $a' \neq a$; or
- (ii) $\psi(a, \pi, \Upsilon(\pi)) < 1 - \varepsilon$, and $\psi(a', \pi, \Upsilon(\pi)) \geq \varepsilon/A$ and $|p(a', \pi, \Upsilon(\pi)) - \pi| > \varepsilon$ at some a' .

Proof. Choose $\eta > 0$ so small that $\psi(a', \pi, \Upsilon(\pi)) < 1 - \eta$ for any action $a' \neq a$ and all π within η of action a ’s cascade set $C_a(\delta)$. If such η does not exist, a.s. taking some action a' is optimal at some $\tilde{\pi} \in C_a(\delta)$, since the optimal rule correspondence is u.h.c. This is impossible: always taking a' incurs a strict myopic loss and no information gain.

CASE 1: BOUNDED SIGNALS. By Claim 5(b), for π close enough to 0 or 1, active learning optimally stops. So, we need only consider π in some closed subinterval $I \subset (0, 1)$. As some signals are informative, $\min \text{supp}(F) \equiv \sigma_0 < 1/2 < \sigma_1 \equiv \max \text{supp}(F)$. Define $S = \text{supp}(F) \setminus ((2\sigma_0 + 1)/4, (2\sigma_1 + 1)/4)$. Let $\eta_1 = \min_{\pi \in I, \sigma \in S} |R(\sigma, \pi) - \pi| > 0$, and

$$\eta_2 = \min_{\pi \in I} \left\{ \left| \frac{\pi F^H(\frac{1}{2})}{\pi F^H(\frac{1}{2}) + (1 - \pi) F^L(\frac{1}{2})} - \pi \right|, \left| \frac{\pi(1 - F^H(\frac{1}{2}))}{\pi(1 - F^H(\frac{1}{2})) + (1 - \pi)(1 - F^L(\frac{1}{2}))} - \pi \right| \right\} > 0.$$

Choose $\varepsilon = \min\{\eta, \eta_1, \eta_2, F^H((2\sigma_0 + 1)/4), 1 - F^L((2\sigma_1 + 1)/4)\}$.

ITEM (I): When $\psi(a, \pi, \Upsilon(\pi)) \geq 1 - \varepsilon$, by Corollary 1, any $a' \neq a$ is only taken for $\sigma \in S$. Since $p(a', \pi, \Upsilon(\pi))$ averages $R(\sigma, \pi)$ over σ to a' , $|p(a', \pi, \Upsilon(\pi)) - \pi| \geq \eta_1 \geq \varepsilon$.

ITEM (II): Suppose $\psi(a, \pi, \Upsilon(\pi)) < 1 - \varepsilon$. Any $a' \neq a$ likewise has $\psi(a', \pi, \Upsilon(\pi)) < 1 - \eta \leq 1 - \varepsilon$. Then there exists a'' with $\psi(a'', \pi, \Upsilon(\pi)) > \varepsilon/A$ which is not taken at $\sigma = 1/2$. Then $|p(a'', \pi, \Upsilon(\pi)) - \pi| \geq \eta_2 \geq \varepsilon$.

CASE 2: UNBOUNDED SIGNALS. By Lemma 1, v has absolute slope at most $\kappa = \max\{|u_H(1) - u_L(1)|, |u_H(A) - u_L(A)|\} < \infty$. Since action 1 (A) is strictly optimal in state L (H), there exists $\eta_3 \in (0, 1)$ such that: for all $a \neq 1$, $(1 - \delta)(\bar{u}(1, r) - \bar{u}(a, r)) > A\delta\kappa\eta_3$ when $r \in [0, \eta_3]$, and, for all $a \neq A$, $(1 - \delta)(\bar{u}(A, r) - \bar{u}(a, r)) > A\delta\kappa\eta_3$ when $r \in [1 - \eta_3, \eta_3]$. Let U denote the maximal possible myopic payoff difference among any pair of actions. Choose $\varepsilon = \min\{\eta, \eta_3/2, (A - 1)\delta\kappa\eta_3/[2A\delta\kappa\eta_3 + 2(1 - \delta)U/A]\}$. WLOG, focus on action 1.

ITEM (I): Bayes rule gives $p(1, \pi, \Upsilon(\pi)) \leq \pi/(1 - \varepsilon) < \eta_3$. Towards a contradiction, suppose active $a \neq 1$ has $|p(a, \pi, \Upsilon(\pi)) - \pi| \leq \varepsilon$, so $p(a, \pi, \Upsilon(\pi)) < \eta_3$. Now, the planner gains by merging actions a and 1, directing all these signals to 1. The continuation belief remains in $[0, \eta_3]$, so the future value loss is at most $\psi(a, \pi, \Upsilon(\pi))\delta 2\kappa\eta_3$. The myopic value gain is at least $\psi(a, \pi, \Upsilon(\pi))(1 - \delta)(\bar{u}(1, p(a, \pi, \Upsilon(\pi)))) - \bar{u}(a, \pi, \Upsilon(\pi)) > \psi(a, \pi, \Upsilon(\pi))A\delta\kappa\eta_3$. This contradiction to optimality of $\Upsilon(\pi)$ proves the desired $|p(a, \pi, \Upsilon(\pi)) - \pi| > \varepsilon$.

ITEM (II): Towards a contradiction, suppose all actions with $\psi(a, \pi, \Upsilon(\pi)) > \varepsilon/A$ have $|p(a, \pi, \Upsilon(\pi)) - \pi| \leq \varepsilon$. Then $\tilde{A} = \{a | p(a, \pi, \Upsilon(\pi)) < \eta_3\}$ has chance $\psi(\tilde{A}, \pi, \Upsilon(\pi)) \geq 1 - \varepsilon$. The proof of item (i) shows that $1 \notin \tilde{A}$. Now another gain is available: Take 1 where any action in \tilde{A} was taken, and take the arbitrary $\tilde{a} \in \tilde{A}$ where 1 was taken. The future value loss from pooling \tilde{A} is at most $\delta\kappa\eta_3$, with chance $\psi(\tilde{A}, \pi, \Upsilon(\pi)) < 1$. Since 1 replaces worse actions, there is a myopic gain of at least $A\delta\kappa\eta_3$ with chance $\psi(\tilde{A}, \pi, \Upsilon(\pi)) \geq 1 - \varepsilon$. There is possibly a myopic loss of at most $(1 - \delta)U$ with chance $\psi(1, \pi, \Upsilon(\pi)) \leq \varepsilon/A$. This sums to an overall gain since $(1 - \varepsilon)A\delta\kappa\eta_3 - \varepsilon(1 - \delta)U/A > \delta\kappa\eta_3$ by choice of ε . \parallel

20. This Lemma is Corollary 5.29 of Breiman (1968): Let event \mathcal{A}_n be measurable w.r.t. (Y_1, \dots, Y_n) for a stochastic process Y_1, Y_2, \dots . Then, a.s. $\{\mathcal{A}_n \text{ infinitely often}\} = \{\sum_{n=1}^{\infty} P(\mathcal{A}_{n+1} | Y_n, \dots, Y_1) = \infty\}$.

D. IMPLEMENTATION PROOFS

PIVOT MECHANISM: PROOF OF PROPOSITION 3. In a cascade on a , $t=0$ solves (6.1) and (6.2), as selfishness is optimal ($C_a(\delta) \subseteq C_a(0)$).

CASE 1. $\pi \notin C(\delta)$, AND THAT NO ACTIVE ACTION HITS THE CASCADE SET. Consider any active action a . Since the next agent is not in a cascade set, we have $F^L(\hat{\sigma}) > F^H(\hat{\sigma})$. Then there exists unique transfers $t(a, \hat{a})$ and $t(a, -\hat{a})$ solving (6.2), since the determinant is $F^L(\hat{\sigma})(1 - F^H(\hat{\sigma})) - (1 - F^L(\hat{\sigma}))F^H(\hat{\sigma}) > 0$. These transfers deliver the right incentive: The agent with posterior belief r expects to receive $r[F^H(\hat{\sigma})t(a, \hat{a}) + (1 - F^H(\hat{\sigma}))t(a, -\hat{a})] + (1 - r)[F^L(\hat{\sigma})t(a, \hat{a}) + (1 - F^L(\hat{\sigma}))t(a, -\hat{a})] = rt_H(a|\pi) + (1 - r)t_L(a|\pi)$ from action a . Then $\tilde{u}(a, r) + rt_H(a|\pi) + (1 - r)t_L(a|\pi)$ is an affine transformation of the index $w(a, \pi, r)$ in (4.3), where the transformation depends on π and not r . All told, this solves (6.1).

CASE 2. $\pi \notin C(\delta)$, AND A CONTINUATION IS IN A CASCADE. If action a sparks a cascade, equations (6.2) might not be solvable, as $F^L(\hat{\sigma}) = F^H(\hat{\sigma})$ (i.e. zero or one). But by Claim 7, at most one cascade set, say $C_{a'}(\delta)$, can possibly be hit by all actions. So inspired, we devise a non-pivot mechanism: choose a zero transfer for active actions leading to the cascade set $C_{a'}(\delta)$. For other active actions a , the transfer pays the difference between continuation values of a and a' . Since $v_\omega(p(a', \pi, \xi)) = u_\omega(a')$, transfers solve (6.2), where $t'_\omega(a|\pi) = [\delta/(1 - \delta)](v_\omega(p(a, \pi, \xi)) - u_\omega(a'))$. Also, $\tilde{u}(a, r) + rt'_H(a|\pi) + (1 - r)t'_L(a|\pi)$ is again an affine transformation of the index $w(a, \pi, r)$ in (4.3), independent of r .

We can easily deter agents from taking inactive actions with large negative transfers. \parallel

MIMICRY WITH TWO ACTIONS: PROOF OF COROLLARY 3. Assume $a = 1$, with $p(1, \pi, \xi) \notin C(\delta)$. Assume first the case where $p(2, \pi, \xi) \notin C(\delta)$. By (6.2),

$$t(1, 1) - t(1, 2) = \frac{t_L(1|\pi) - t_H(1|\pi)}{\psi(1, L, \xi) - \psi(1, H, \xi)}. \quad (\text{D.1})$$

Since the two actions are taken in the myopic order, $\psi(1, L, \xi) > \psi(1, H, \xi)$. Thus, the fraction has the sign of the numerator. The definition of t in (6.1) and $v_H(p) - v_L(p) = \tau(p, 1) - \tau(p, 0) = v'(p)$ imply $t_L(1|\pi) - t_H(1|\pi) = [\delta/(1 - \delta)](v'(\pi) - v'(p(1, \pi, \xi)))$. By the myopic action ordering, we have $p(1, \pi, \xi) < \pi$. By Claim 4, the value function is strictly convex and thus $v'(\pi) - v'(p(1, \pi, \xi)) > 0$. Thus $t(1, 1) - t(1, 2) > 0$. When $p(2, \pi, \xi) \in C(\delta)$, the logic is the same, substituting t in (D.1) by t' . \parallel

E. POSTERIOR MONOTONICITY PROOFS

Claim 14 Given (PM), private signals have a continuous cdf, increasing on an interval.

Proof. Assume a non-convex private signal support. Pick any $\sigma_1 \in \text{co}(\text{supp}(F)) \setminus \text{supp}(F)$. Let σ_0 be the upper bound of $\text{supp}(F) \cap [0, \sigma_1]$, and σ_2 the lower bound of $\text{supp}(F) \cap (\sigma_1, 1]$. Pick payoffs with $I_a = [\sigma_0, \sigma_2]$ the posterior belief interval for some action a (Corollary 1). By (3.2), the posterior map $R(\pi, \sigma)$ is continuous and monotone, with $R(1/2, \sigma) \equiv \sigma$. Fix π near $1/2$ but $\pi < 1/2$. Then, a positive probability $\psi(a, \pi) > 0$ of private signals above σ_2 map to a posterior in I_a . Since (σ_1, σ_2) is a missing interval in $\text{supp}(F)$, no private signals below σ_2 map to a , and $\sigma_2 > \sigma_1$, and $R(\pi, \sigma)$ is continuous, the continuation public belief $p(a, \pi)$ (an average of these private signals) exceeds σ_1 . Similarly, for π' near $1/2$ and $\pi' > 1/2$, the continuation public belief $p(a, \pi')$ is below σ_1 . This contradicts (PM).

Assume next, for a contradiction a positive mass of private signals at σ_1 . As there are no perfectly revealing signals, $\sigma_1 \in (0, 1)$. Since some signals are informative, $\text{supp}(F)$ is not a single-point. So $\text{supp}(F)$ contains $[\sigma_1, \sigma_1 + \eta]$ or $[\sigma_1 - \eta, \sigma_1]$, for some $\eta > 0$. In the first case,

$$\hat{\sigma}_0 \equiv E[\sigma | \sigma \in [\sigma_1, \sigma_1 + \eta]] < E[\sigma | \sigma \in (\sigma_1, \sigma_1 + \eta)] \equiv \hat{\sigma}_2,$$

since σ_1 has positive probability. Assume payoffs so that action b has the posterior interval $I_b = [\sigma_1, \sigma_1 + \eta]$. Since σ_1 maps to I_b , we have $\psi(b, \pi) > 0$ if $\pi < 1/2$ is near $\frac{1}{2}$, and $\psi(b, \pi') > 0$ if $\pi' > 1/2$ is near $\frac{1}{2}$. (In both cases, the posterior is near the private signal σ_1 .) As the posterior $R(\pi, \sigma)$ is continuous, the continuation public belief $p(b, \pi)$ tends to a limit at least $R(1/2, \hat{\sigma}_2)$, as π approaches $1/2$ from below, and at most $R(1/2, \hat{\sigma}_0)$, as π approaches $1/2$ from above. Since $\hat{\sigma}_0 < \hat{\sigma}_2$, this contradicts (PM). If the support is $[\sigma_1 - \eta, \sigma_1]$, the posterior interval $[\sigma_1 - \eta, \sigma_1]$ likewise provides a non-monotonicity at prior $1/2$. \parallel

PROOF OF PROPOSITION 4. The private signal density induces a density $g(r|\pi)$ on posterior beliefs $r = R(\pi, \sigma)$ on state H . As in Section 7, write the signal $\sigma = \mathcal{S}(\ell) = e^\ell / (1 + e^\ell)$ in terms of the log likelihood ratio $\ell = \Lambda(\sigma) = \log(\sigma / (1 - \sigma))$, with state- ω contingent density $\phi^\omega(\ell) \equiv f^\omega(\mathcal{S}(\ell))\mathcal{S}'(\ell)$. Denote by $\phi(\ell|\pi) = \pi\phi^H(\ell) + (1 - \pi)\phi^L(\ell)$ the unconditional density for the log-likelihood ratio ℓ of the private signal when the public belief is π .

By Bayes rule, the posterior log-likelihood ratio is $\rho \equiv \Lambda(r) = \ell + \Lambda(\pi)$, and thus has density $\phi(\rho - \Lambda(\pi)|\pi)$. By changing variable from r to ρ , the continuation public belief in §4 is

$$p(a, \pi) = \frac{\int_{I_a} rg(r|\pi)dr}{\int_{I_a} g(r|\pi)dr} = \frac{\int_{\Lambda(I_a)} \mathcal{S}(\rho)\phi(\rho - \Lambda(\pi)|\pi)d\rho}{\int_{\Lambda(I_a)} \phi(\rho - \Lambda(\pi)|\pi)d\rho}, \quad (\text{E.1})$$

writing $p(a, \pi)$ for $p(a, \pi, \xi)$. As $\mathcal{S}' > 0$, (PM) ensues if ϕ is *log-supermodular* (LSPM).

Step 1 (An Equivalence) *The density ϕ^L is strictly logconcave iff $\phi(\rho - \Lambda(\pi)|\pi)$ is strictly log-supermodular in (ρ, π) wherever $\rho - \Lambda(\pi)$ is in the support $\text{supp}(\phi)$.*

Proof. Since $\phi^H(\ell) = e^\ell \phi^L(\ell)$ (see §7), we have $\phi(\rho - \Lambda(\pi)|\pi) = (1 - \pi)(1 + e^\rho) \phi^L(\rho - \Lambda(\pi))$, since $\pi e^{-\Lambda(\pi)} \equiv 1 - \pi$. Strict LSPM of ϕ is equivalent to strict LSPM of $\phi^L(\rho - \Lambda(\pi))$. As Λ is strictly increasing, this holds iff ϕ^L is strictly log-concave, as Karlin (1968) shows. \parallel

Step 2 (Necessity) *$\phi(\rho - \Lambda(\pi)|\pi)$ strictly LSPM in (ρ, π) if $\rho - \Lambda(\pi) \in \text{supp}(\phi) \Rightarrow$ (PM).*

Proof. Fix $\pi' > \pi$ and an active action a (so $\psi(a, \pi), \psi(a, \pi') > 0$). Activity implies that $\Lambda(I_a) - \Lambda(\pi)$ and $\Lambda(I_a) - \Lambda(\pi')$ overlap $\text{supp}(\phi)$. Next, if $\rho' > \rho$, and $\rho' - \Lambda(\pi') > \rho - \Lambda(\pi')$ are in the convex $\text{supp}(\phi)$, then so are $\rho - \Lambda(\pi)$ and $\rho' - \Lambda(\pi')$, as Λ increases. As $\phi(\rho - \Lambda(\pi)|\pi)$ is strictly LSPM in (ρ, π) , $\phi(\rho - \Lambda(\pi)|\pi) \phi(\rho' - \Lambda(\pi')|\pi') > \phi(\rho' - \Lambda(\pi)|\pi) \phi(\rho - \Lambda(\pi')|\pi')$, lest the left side vanishes. So (E.1) strictly increases in π , by Karlin and Rubin (1956). \parallel

Step 3 (Sufficiency) *Suppose $\phi(\rho - \Lambda(\pi)|\pi)$ is positive and continuous in ρ on its support, for every π . If (PM) holds, then $\phi(\rho - \Lambda(\pi)|\pi)$ is strictly LSPM when $\rho - \Lambda(\pi) \in \text{supp}(\phi)$.*

Proof. If LSPM fails, there is $\pi_2 > \pi_1$ and $\rho_2 > \rho_1$ with $\rho_1 - \Lambda(\pi_1)$ and $\rho_2 - \Lambda(\pi_2)$ in $\text{supp}(\phi)$ and (\diamond) : $\phi(\rho_1 - \Lambda(\pi_1)|\pi_1) \phi(\rho_2 - \Lambda(\pi_2)|\pi_2) < \phi(\rho_2 - \Lambda(\pi_1)|\pi_1) \phi(\rho_1 - \Lambda(\pi_2)|\pi_2)$. Define

$$H(x) \equiv [\log \phi(\rho_1 - \Lambda(\pi_2)|\pi_2) - \log \phi(\rho_1 - \Lambda(\pi_1)|\pi_1)] - [\log \phi(x - \Lambda(\pi_2)|\pi_2) - \log \phi(x - \Lambda(\pi_1)|\pi_1)].$$

Then, H is continuous on $[\rho_1, \rho_2]$, with $H(\rho_1) = 0 < H(\rho_2)$ by (\diamond) . Let $x_2 = \max\{x \in [\rho_1, \rho_2] : H(x) = \frac{1}{2}H(\rho_2)\}$, where $\rho_1 < x_2 < \rho_2$. By definition of x_2 , if $x \in (x_2, \rho_2)$, then $H(x) > H(x_2)$, and so (by rewriting) $\phi(x - \Lambda(\pi_2)|\pi_2) / \phi(x_2 - \Lambda(\pi_2)|\pi_2) < \phi(x - \Lambda(\pi_1)|\pi_1) / \phi(x_2 - \Lambda(\pi_1)|\pi_1)$. Integrating this over $x \in (x_2, \rho_2)$, and using continuity near x_2 , there exists $\varepsilon \in (0, x_2 - \rho_1)$ such that

$$\frac{\int_{x_2}^{\rho_2} \phi(x - \Lambda(\pi_2)|\pi_2) dx}{\int_{x_2-\varepsilon}^{\rho_2} \phi(x - \Lambda(\pi_2)|\pi_2) dx} < \frac{\int_{x_2}^{\rho_2} \phi(x - \Lambda(\pi_1)|\pi_1) dx}{\int_{x_2-\varepsilon}^{\rho_2} \phi(x - \Lambda(\pi_1)|\pi_1) dx}. \quad (\text{E.2})$$

Define the cdf G^i on $[x_2 - \varepsilon, \rho_2]$ by $G^i(x) = \int_{x_2-\varepsilon}^x \phi(t - \Lambda(\pi_i)|\pi_i) dt / \int_{x_2-\varepsilon}^{\rho_2} \phi(t - \Lambda(\pi_i)|\pi_i) dt$. Since $G^1(x_2 - \varepsilon) = G^2(x_2 - \varepsilon) = 0$ and $G^1(\rho_2) = G^2(\rho_2) = 1$, let $\underline{x} = \max\{x \in [x_2 - \varepsilon, x_2] : G^1(x) = G^2(x)\}$ and $\bar{x} = \min\{x \in [x_2, \rho_2] : G^1(x) = G^2(x)\}$. Then G^1 strictly stochastically dominates G^2 on $[\underline{x}, \bar{x}]$, by (E.2). If action a is taken for $\rho \in [\underline{x}, \bar{x}]$, we contradict (PM):

$$p(a, \pi_2) = \frac{\int_{\underline{x}}^{\bar{x}} \mathcal{S}(\rho_1) \phi(\rho_1 - \Lambda(\pi_2)|\pi_2) d\rho_1}{\int_{\underline{x}}^{\bar{x}} \phi(\rho_1 - \Lambda(\pi_2)|\pi_2) d\rho_1} < \frac{\int_{\underline{x}}^{\bar{x}} \mathcal{S}(\rho_1) \phi(\rho_1 - \Lambda(\pi_1)|\pi_1) d\rho_1}{\int_{\underline{x}}^{\bar{x}} \phi(\rho_1 - \Lambda(\pi_1)|\pi_1) d\rho_1} = p(a, \pi_1).$$

Next assume strict LSPM fails. So equality holds in (\diamond) for some $\pi_2 > \pi_1$ and $\rho_2 > \rho_1$, so that $\phi(\rho_1 - \Lambda(\pi_1)|\pi_1) \phi(\rho_2 - \Lambda(\pi_2)|\pi_2) = \phi(\rho_2 - \Lambda(\pi_1)|\pi_1) \phi(\rho_1 - \Lambda(\pi_2)|\pi_2)$. Now, $H(x) \not\equiv 0$ on $[\rho_1, \rho_2]$, for otherwise $p(a, \pi_2) = p(a, \pi_1)$ when a is taken on $[\underline{x}, \bar{x}]$, a contradiction to (PM). So, we must either have $H(x') < 0$ or $H(x') > 0$ for some $x' \in (\rho_1, \rho_2)$. This respectively contradicts (the just proved) LSPM on $\{x', \rho_2\} \times \{\pi_1, \pi_2\}$ or $\{\rho_1, x'\} \times \{\pi_1, \pi_2\}$. \parallel

Step 4 *If (PM) holds, then $\phi(\rho - \Lambda(\pi)|\pi)$ is strictly LSPM when $\rho - \Lambda(\pi) \in \text{supp}(\phi)$.*

Proof. Assuming (PM), we extend Step 3. Let (X_m) be a sequence of mean-0 normal r.v.s with vanishing variance, and independent of ρ, θ . Let ρ_m denote the posterior belief for someone who observes $\rho + X_m$, and let $\phi_m(\rho_m|\pi)$ denote its conditional density. Since the density Υ_m of X_m is log-concave, the pair (ρ, ρ_m) satisfies the MLRP. In particular, for any given X_m , posterior ρ_m is a continuous, increasing function of ρ .

We show that the distribution of ρ_m inherits property (PM) for all m . Let thus any interval J be given. For any realization X_m , let $J_m(X_m)$ denote the interval such that $\rho_m \in J$ iff $\rho \in J_m(X_m)$. The posterior belief given $\{\rho_m \in J\}$ is $\int_{-\infty}^{\infty} p(J_m(x), \pi) \Upsilon_m(x) dx$. By (PM), for every x , $p(J_m(x), \pi)$ is weakly increasing in π , and strictly so wherever $J_m(x)$ intersects the support of ρ . The integral thus inherits (PM), as desired.

Since X_m and thus ρ_m have distributions that satisfy the auxiliary assumptions of Step 3, we can conclude that $\phi_m(\rho_m - \Lambda(\pi)|\pi)$ is strictly LSPM. Step 1 implies that ϕ_m is strictly logconcave. As m tends to infinity, the variance of X_m vanishes, and in the limit ϕ is log-concave. If ϕ is log-linear on any interval, π and J can be chosen so $p(J, \pi)$ is locally constant in π , in contradiction to (PM). \parallel

F. CONTRARIANISM PROOFS

Lemma 2 (Tangents to a Convex Function) Assume consecutive tangents $\tau_i < \tau_{ii} < \tau_{iii}$ to a value function v at $z_i < z_{ii} < z_{iii}$. Then $\tau_{ii}(z_i) \geq \tau_{iii}(z_i)$ (respectively, $\tau_i(z_{iii}) \leq \tau_{ii}(z_{iii})$), with strict inequality unless v is affine on $[z_{ii}, z_{iii}]$ (respectively, on $[z_i, z_{ii}]$).

Proof. When v is affine on $[z_i, z_{ii}]$, subtangents τ_i and τ_{ii} can coincide, with $\tau_i(z_{iii}) = \tau_{ii}(z_{iii})$. Otherwise, τ_{ii} is steeper than τ_i . Thus, $\tau_{ii}(z_{iii}) - \tau_{ii}(z_{ii}) > \tau_i(z_{iii}) - \tau_i(z_{ii})$, whence $\tau_{ii}(z_{iii}) - \tau_i(z_{iii}) > \tau_{ii}(z_{ii}) - \tau_i(z_{ii})$. Since v is convex, the subgradient τ_i lies below v at z_{ii} , so that $\tau_{ii}(z_{ii}) = v(z_{ii}) \geq \tau_i(z_{ii})$. So $\tau_{ii}(z_{iii}) > \tau_i(z_{iii})$. The z_i analysis is similar. \parallel

By assumption, at public beliefs $\pi < \pi'$, there exist optima with the same action order. The optimal rules at π and π' therefore also solve the Bellman problem (4.1) with (4.2) when we restrict the choice set to this action order. In this restricted problem, we explore the comparative statics properties of the constrained Bellman equation for any belief outside the cascade set $C(\delta)$. Define the constrained Bellman function as the right side of (4.2):

$$B(\theta|\pi) = \sum_{a=1}^A \psi(a, \pi, \theta) [(1-\delta)\tilde{u}(a, p(a, \pi, \theta)) + \delta v(p(a, \pi, \theta))]. \quad (\text{F.1})$$

Solutions to the constrained problem $\max_{\theta \in \Theta(\pi)} B(\theta|\pi)$ define an optimizer set $\Theta^*(\pi)$. To prove Proposition 5, it suffices that $\Theta^*(\pi)$ increase in the strong set order.

F.1. Proof of Proposition 5 with two actions

We wish to apply a clever comparative statics result in Quah and Strulovici (2009). Their Theorem 1 delivers our conclusion provided $B(\cdot|\pi')$ exceeds $B(\cdot|\pi)$ in their interval dominance order. A sufficient condition for this order is their Proposition 2, that there exist an increasing and strictly positive function $\alpha(\theta)$ with $B_\theta(\theta|\pi') \geq \alpha(\theta)B_\theta(\theta|\pi)$. Inspired by (4.3) and (F.1), we derive an expression for $B_\theta(\theta|\pi)$ in terms of the welfare index.

Lemma 3 (FOC) The Bellman function B is differentiable almost everywhere with derivative

$$B_\theta(\theta|\pi) = g(\theta, \pi) (w(1, \pi, \theta) - w(2, \pi, \theta)). \quad (\text{F.2})$$

Also, B is absolutely continuous, with $B(\theta'|\pi) - B(\theta|\pi) = \int_\theta^{\theta'} B_\theta(\tilde{\theta}|\pi) d\tilde{\theta}$ for $\theta, \theta' \in \Theta(\pi)$.

Proof. From (F.1), the Bellman function is a.e. differentiable in θ . For by assumption (LC), $p(a, \pi, \theta)$ is strictly monotone and differentiable, and the convex function v is differentiable a.e. Since \tilde{u} and τ_a are affine functions, and since $p(a, \pi, \xi) = \int_{\xi-1(a)} r(\pi, \sigma) dF^\pi$, we can use Proposition 1 to rewrite (4.2) as follows, proving Lemma 3:

$$B(\theta|\pi) = \int_0^\theta w(1, \pi, r) g(r|\pi) dr + \int_\theta^1 w(2, \pi, r) g(r|\pi) dr. \quad (\text{F.3})$$

\parallel

Returning to the proof of Proposition 5, suppose that the thresholds $\theta \in \Theta^*(\pi)$ and $\theta' \in \Theta^*(\pi')$ are inversely ordered as $\theta' < \theta$ —otherwise, we're done. Since $r(\sigma, \pi)$ increases in π , the open interval $\Theta(\pi)$ rises in π . So $[\theta', \theta] \subset \Theta(\pi) \cap \Theta(\pi')$. We first argue that the index difference $\Delta(\tilde{\theta}, \pi) \equiv w(1, \pi, \tilde{\theta}) - w(2, \pi, \tilde{\theta})$ in (F.2) weakly increases in the public belief π , when $\tilde{\theta} \in [\theta', \theta]$. By Proposition 4, continuation beliefs rise in public beliefs: $p(a, \pi', \tilde{\theta}) > p(a, \pi, \tilde{\theta})$ for $a = 1, 2$. Using definition (4.3), Lemma 2 yields the desired,

$$\Delta(\tilde{\theta}, \pi') - \Delta(\tilde{\theta}, \pi) = \delta [\tau_1'(\tilde{\theta}) - \tau_1(\tilde{\theta})] + [\tau_2(\tilde{\theta}) - \tau_2'(\tilde{\theta})] \geq 0. \quad (\text{F.4})$$

Next, $\alpha(\tilde{\theta}) \equiv g(\tilde{\theta}|\pi')/g(\tilde{\theta}|\pi)$ is a positive and non-decreasing function over $[\theta', \theta]$, since g is log-supermodular, by Lemma 4. Then Lemma 3 and inequality (F.4) imply:

$$B_\theta(\tilde{\theta}|\pi') = g(\tilde{\theta}|\pi') \Delta(\tilde{\theta}, \pi') \geq g(\tilde{\theta}|\pi) \Delta(\tilde{\theta}, \pi) = \alpha(\tilde{\theta}) B_\theta(\tilde{\theta}|\pi), \quad (\text{F.5})$$

This implies that B obeys the interval dominance order, by Proposition 2 in Quah and Strulovici (2009). By their Theorem 1, $\Theta(\pi)$ rises in the strong set order—contrarianism.

Consider the stronger claim in Proposition 5 that the optimizer set *strictly rises*. Suppose first that thresholds $\theta \geq \theta'$ are respectively optimal at public beliefs $\pi < \pi'$. By the already proven strong set order, $\theta \in \Theta^*(\pi')$. By Proposition 1, $w(1, \pi, \theta) - w(2, \pi, \theta) = w(1, \pi', \theta) - w(2, \pi', \theta) = w(1, \pi', \theta') - w(2, \pi', \theta') = 0$. The first difference vanishes since θ is optimal at π , the second since θ is optimal at π' , and the third since θ' is optimal at π' . If $\theta > \theta'$, we contradict the fact that $w(2, \pi', r) - w(1, \pi', r)$ increases in r , as follows from (4.3). For the natural action order implies that $\tilde{u}(2, r) - \tilde{u}(1, r)$ is strictly increasing, and convexity of v implies that its tangent difference $\tau_2'(r) - \tau_1'(r)$ is monotone.

Consider the other possibility with $\theta = \theta'$. Now $\pi < \pi'$ implies $p(a, \pi, \theta) < p(a, \pi', \theta)$, and at least one of $p(1, \pi', \theta), p(2, \pi, \theta)$ is outside the cascade set, by Claim 7. Lemma 2 gives the contradiction $w(1, \pi, \theta) - w(2, \pi, \theta) > w(1, \pi', \theta) - w(2, \pi', \theta)$. The inequality is strict because v is strictly convex outside the cascade set, by Claim 4. \parallel

F.2. Proof of Proposition 5 for multiple actions

Recall from the proof of Claim 7 that $\text{supp}(F) = [\underline{\sigma}, \bar{\sigma}]$.

Claim 15 Let $\theta \in \Theta(\pi)$ obey $\theta_0 = r(\underline{\sigma}, \pi)$ and $\theta_A = r(\bar{\sigma}, \pi)$. Assume $\theta_a = \dots = \theta_{a+j} = x$ for some $a \geq 1$ and $j \geq 0$ with $a+j \leq A-1$, and $\theta_{a-1} < x < \theta_{a+j+1}$. Then the Bellman function B in (F.7) is absolutely continuous in x , and its derivative in x almost everywhere equals:

$$B_x(\theta|\pi) \equiv g(x|\pi)(w(a, \pi, x) - w(a+j+1, \pi, x)). \quad (\text{F.6})$$

Also, for all $\pi'' > \pi'$, there exists a positive and increasing function $\alpha(x)$ such that the Bellman function $B(\theta|\pi)$ a.e. obeys $B_x(\theta|\pi'') \geq \alpha(x)B_x(\theta|\pi')$ when $\theta \in \Theta(\pi') \cap \Theta(\pi'')$.

The proof of this many action generalization follows closely on Lemma 3, since we take action a for $r \in [\theta_{a-1}, x]$, and action $a+j+1$ for $r \in [x, \theta_{a+j+1}]$. So the derivative of the Bellman function B in x is similar to (F.2) which had payoffs and tangents for actions $a=1$ and $a+j+1=2$. Thus, (F.6) follows. The inequality follows similarly from (F.5). \parallel

For contrarianism, we must show that $\Theta^*(\pi')$ exceeds $\Theta^*(\pi)$ in the strong set order.²¹

Claim 16 The threshold space $\Theta(\pi)$ is a lattice, and B is supermodular for $\theta \in \Theta(\pi)$.

Proof. Assume $\theta, \theta' \in \Theta(\pi)$. Then $\theta \wedge \theta' \in \Theta(\pi)$ since $(\theta \wedge \theta')_a = \theta_a \wedge \theta'_a \leq \theta_{a+1} \wedge \theta'_{a+1} = (\theta \wedge \theta')_{a+1}$ for every a . Similarly, $\theta \vee \theta' \in \Theta(\pi)$. Next, to show that B is supermodular in θ , let $\theta'_a > \theta_a$. If θ_{-a} increases, both continuation beliefs $p(a, \pi, \theta)$ and $p(a+1, \pi, \theta)$ increase. Since $p(a, \pi, \theta) < \theta_a < p(a+1, \pi, \theta)$, Lemma 2 implies that $w(a, \pi, \theta_a)$ increases while $w(a+1, \pi, \theta_a)$ decreases. So the difference $w(a, \pi, \theta_a) - w(a+1, \pi, \theta_a)$ increases in θ_{-a} . Then by (F.6), the Bellman difference $B(\theta'_a, \theta_{-a}) - B(\theta_a, \theta_{-a})$ increases in θ_{-a} . Supermodularity can now be decomposed into a summation of differences of this form. \parallel

Fixing the action ordering, the Bellman function (4.2) for a convex continuation value v is:

$$B(\theta|\pi) = \sum_{a=1}^A \psi(a, \pi, \theta) [(1-\delta)\bar{u}(a, p(a, \pi, \theta)) + \delta v(p(a, \pi, \theta))]. \quad (\text{F.7})$$

We now prove Proposition 5 for finitely many actions. Pick beliefs $\pi < \pi'$ and assume that $\theta \in \Theta^*(\pi)$ and $\theta' \in \Theta^*(\pi')$. If $\theta \leq \theta'$, we are done. Assume next that they are inversely ordered $\theta' < \theta$. We verify $\theta \in \Theta^*(\pi')$ and $\theta' \in \Theta^*(\pi)$. First, both $[\theta_1, \theta_{A-1}]$ and $[\theta'_1, \theta'_{A-1}]$ are subsets of $\Theta(\pi) \cap \Theta(\pi')$, since $[\theta_1, \theta_{A-1}] \subset \Theta(\pi)$ and $[\theta'_1, \theta'_{A-1}] \subset \Theta(\pi')$ and $[\theta_1, \theta_{A-1}]$ lies above $[\theta'_1, \theta'_{A-1}]$ in the strong set order, and yet $\Theta(\pi)$ lies below $\Theta(\pi')$ in the strong set order. Second, let X be the set of all cut-off rules with cut-off points in $\Theta(\pi) \cap \Theta(\pi')$. By Tian (2014), $B(\cdot|\pi')$ dominates $B(\cdot|\pi)$ in the interval dominance order over X since, by Claim 15, the condition for Proposition 2 in Tian (2014) is satisfied.

Finally, suppose that θ and θ' are not ordered. We now need a stronger proof ingredient—specifically, we exploit the supermodularity of B (Claim 16). Our result follows if:

$$B(\theta|\pi) - B(\theta \wedge \theta'|\pi) \geq 0 (>0) \implies B(\theta \vee \theta'|\pi') - B(\theta'|\pi') \geq 0 (>0). \quad (\text{F.8})$$

Let's see why this suffices. Since θ is optimal at π , the left side is non-negative, and thus $\theta \vee \theta'$ is optimal at π' by the weak inequality in (F.8). Conversely, if $\theta \wedge \theta'$ is not optimal at π , then θ' is not optimal at π' , by the strict inequality in (F.8).

We split the proof of (F.8) into two parts, since the choice domain $\Theta(\cdot)$ depends on the public belief. Let $(\theta_a, \dots, \theta_{A-1})$ be the components of θ inside $\Theta(\pi')$, for some $a < A$. Choose $z \in \Theta(\pi')$ with $z < \min\{\theta_a, \theta'_1\}$. Let $\hat{\theta} = (z, \dots, z, \theta_a, \dots, \theta_{A-1})$, where the first $a-1$ components are z . Then $\hat{\theta} \in \Theta(\pi) \cap \Theta(\pi')$, since $\theta_{a-1} < z$ follows from $\theta_{a-1} \notin \Theta(\pi')$.

By supermodularity of $B(\cdot|\pi')$, and because $\hat{\theta} \vee \theta' = \theta \vee \theta'$, we have:

$$B(\hat{\theta}|\pi') - B(\hat{\theta} \wedge \theta'|\pi') \geq (>0) \implies B(\theta \vee \theta'|\pi') - B(\theta'|\pi') \geq (>0). \quad (\text{F.9})$$

Then (F.8) follows if we also argue:

$$B(\theta|\pi) - B(\theta \wedge \theta'|\pi) \geq (>0) \implies B(\hat{\theta}|\pi') - B(\hat{\theta} \wedge \theta'|\pi') \geq (>0). \quad (\text{F.10})$$

We now prove (F.10). First, for all $\theta'' \in [\hat{\theta} \wedge \theta', \hat{\theta}]$, we have $\hat{\theta} = \theta \vee \theta''$ and so:

$$B(\hat{\theta}|\pi) - B(\theta''|\pi) \geq B(\theta|\pi) - B(\theta \wedge \theta''|\pi) \geq 0, \quad (\text{F.11})$$

by supermodularity of $B(\cdot|\pi)$ and optimality of θ at π , respectively. When $\theta'' = \hat{\theta} \wedge \theta'$ in (F.11), we have $B(\hat{\theta}|\pi) - B(\hat{\theta} \wedge \theta'|\pi) \geq B(\theta|\pi) - B(\theta \wedge \theta'|\pi)$, since $\theta \leq \hat{\theta}$. Hence, if $B(\theta|\pi) - B(\theta \wedge \theta'|\pi) > 0$, then $B(\hat{\theta}|\pi) - B(\hat{\theta} \wedge \theta'|\pi) > 0$. Finally, the interval dominance ordering of $B(\cdot|\pi')$ over $B(\cdot|\pi)$ lets us conclude (F.10). \parallel

21. Recall that Y' dominates Y in the strong set order if $y \in Y$ and $y' \in Y' \implies y \vee y' \in Y'$ and $y \wedge y' \in Y$.




FIGURE G.1

Inverted action ordering (Section G.1)

Notes: The optimal posterior belief threshold $\theta(\pi)$.

STRICT CONTRARIANISM. Pick $\pi' > \pi$. Let $\theta \in \Theta^*(\pi)$ and $\theta' \in \Theta^*(\pi')$. Then behaviour is contrarian, by Proposition 5. Suppose for a contradiction that it is not strictly so, and thus $\theta'_k \leq \theta_k$ for some k . By Proposition 5, $\theta \vee \theta'$ is optimal under π' . Since $\theta'_k \leq \theta_k$, we have $(\theta \vee \theta')_k = \theta_k$. Suppose that a_j is the highest active action below a_k , and a_m the least active action above a_k . Then $(\theta \vee \theta')_{j-1} < (\theta \vee \theta')_j = \dots = (\theta \vee \theta')_k = \dots = (\theta \vee \theta')_{m-1} < (\theta \vee \theta')_m$, since θ and θ' have the same active actions in natural order. Our proof for two actions then carries over to this case, by considering a neighbouring pair of active actions. \parallel

G. TWO EXAMPLES

G.1. Actions need not be taken in the natural order

To illustrate a non-natural action order (Lesson 2 in Section 4), consider signal densities $f^H(\sigma) = \sigma f(\sigma)$ and $f^L(\sigma) = (1 - \sigma)f(\sigma)$ on $(0, 4/7)$, where $f(\sigma) = 7^8 \sigma^6 / 4^7$. Let action $a = 1, 2$ have payoff $2a - 3$ in state H and $3 - 2a$ in state L , representing payoffs ± 1 when the action matches/mismatches the state. Choose a high discount factor $\delta = 0.95$.

Figure G.1 depicts the numerically calculated private posterior belief threshold $\theta(\pi)$. For public beliefs $\pi \in (0.3, 0.4) \subset (0, 3/7)$, the optimal action order is reversed: action 1 is taken at high signals σ , and action 2 at low signals σ .

To understand this reversion, consider the alternative of switching the two actions, holding fixed the threshold. This switch yields the same information, as it maintains the same chances for the two continuation beliefs. From (4.1), it gives no planner gain when

$$\begin{aligned} & \psi(2, \pi, \xi)(2p(2, \pi, \xi) - 1) + \psi(1, \pi, \xi)(1 - 2p(1, \pi, \xi)) \\ & > \psi(1, \pi, \xi)(2p(1, \pi, \xi) - 1) + \psi(2, \pi, \xi)(1 - 2p(2, \pi, \xi)). \end{aligned} \quad (\text{G.1})$$

Using Bayes rule, $p(a, \pi, \xi) = \pi \psi(a, H, \xi) / \psi(a, \pi, \xi)$, this inequality holds when $\psi(1, \pi, \xi) - \psi(2, \pi, \xi) > 2\pi(\psi(1, H, \xi) - \psi(2, H, \xi))$. Inequality (G.1) holds at low π , as the example shows, when the reversed order takes action 1 for a relatively large set of high signals. \parallel

Note that in this example, the last agent using his own information may take action 2 and push the public belief into the cascade set for action 1. Agents optimally herding on action 1 thus need not follow the lead of the last agent who used private information.

G.2. Contrarianism can fail without posterior monotonicity

We show by an example that (PM) is necessary for contrarianism in Proposition 5 when the convex value function v can be chosen freely in (F.7). We use a version of the two-period professor–student example with $\delta = 1$ in Section 2 to show the principle. The student has three actions available, while the professor has two actions taken in the natural order. The student gets no private signal. The professor’s signal is described by the conditional density $g(r|\pi)$. By assumption, this




FIGURE G.2

Necessity principle (Section G.2)

Notes: The student's value function for Section G.2.

signal structure violates (PM) for some interval, say $[\hat{\theta}, 1]$. Thus,

$$p' \equiv \frac{\int_{\hat{\theta}}^1 rg(r|\pi')dr}{\int_{\hat{\theta}}^1 g(r|\pi')dr} > \frac{\int_{\hat{\theta}}^1 rg(r|\pi'')dr}{\int_{\hat{\theta}}^1 g(r|\pi'')dr} \equiv p''.$$

By this reversal, $\hat{\theta}$ must lie strictly inside the posterior belief supports at π', π'' , so $p'' > \hat{\theta}$.

Figure G.2 illustrates the convex value function that we construct for the example. First choose an arbitrary $\theta_{23} \in (p'', p')$. For any $\varepsilon > 0$, the convex function $\hat{v}(p|\varepsilon)$ consists of three linear segments $\ell_1, \ell_2, \ell_3(\varepsilon)$. Segments ℓ_1, ℓ_2 intersect at $\hat{\theta}$, while $\ell_2, \ell_3(\varepsilon)$ intersect at θ_{23} . ℓ_2 is steeper than ℓ_1 , and the slope of ℓ_3 is $\varepsilon > 0$ higher than ℓ_2 . The intersection of the extended line segments $\ell_1, \ell_3(\varepsilon)$ is denoted $\theta_{13}(\varepsilon)$.

We will show that when $\varepsilon > 0$ is small enough, $\hat{\theta}$ is the unique optimal threshold at belief π'' , while only the strictly higher $\theta_{13}(\varepsilon)$ and θ_{23} are candidates for optimal thresholds at the lower belief π' . In either case, contrarianism fails.

Observe that the three kink points $\hat{\theta}, \theta_{13}(\varepsilon), \theta_{23}$ describe the only candidates for optimal policies. By construction, they are the only ones that solve for index indifference—given discount factor $\delta = 1$, only the tangents to the value function matter. It remains to check sub-optimality of a cascade policy, whereby the posterior is the prior. But the interior threshold $\hat{\theta}$ gives strictly more than $\hat{v}(\pi|0)$ at $\pi = \pi', \pi''$, due to the kink at $\hat{\theta}$.

Consider belief π' . The first order condition fails at $\hat{\theta}$ for any $\varepsilon > 0$, as the tangent at the upper posterior p' is ℓ_3 . So the optimal posterior cut-offs are among $\theta_{13}(\varepsilon)$ and θ_{23} .

Consider π'' . First, suppose we use the cutoff $\theta_{13}(\varepsilon)$. As $\varepsilon \downarrow 0$, the crossing point $\theta_{13}(\varepsilon)$ converges to $\hat{\theta}$, and the upper continuation belief converges to p'' . In other words, it is eventually below θ_{23} , since $p'' < \theta_{23}$. At that point, the tangents at the continuation beliefs after π'' are ℓ_1 and ℓ_2 . These tangents cross at $\hat{\theta}$, and therefore the FOC fails at $\theta_{13}(\varepsilon)$. Second, suppose we use the cutoff θ_{23} . Since $\theta_{23} \in (p'', p')$, it is strictly inside the posterior belief support. Thus, the upper continuation lies in $(\theta_{23}, 1]$, and the lower one either lies in $[0, \hat{\theta})$ or $[\hat{\theta}, \theta_{23})$. If in $[0, \hat{\theta})$, the tangents at the continuation beliefs are ℓ_1 and $\ell_3(\varepsilon)$. These cross at $\theta_{13}(\varepsilon)$, and so the FOC fails at θ_{23} . If in $[\hat{\theta}, \theta_{23})$, the FOC holds. But as $\varepsilon \downarrow 0$, the continuation value approaches $\hat{v}(\pi''|0)$. But as noted before, $\hat{\theta}$ yields a strictly higher continuation value than $\hat{v}(\pi''|0)$. \parallel

REFERENCES

- AKERLOF, G. A. (1970), "The Market for "Lemons": Quality Uncertainty and the Market Mechanism", *Quarterly Journal of Economics*, **84**, 488–500.
- BANERJEE, A. V. (1992), "A Simple Model of Herd Behavior", *Quarterly Journal of Economics*, **107**, 797–817.
- BERGEMANN, D. and VÄLIMÄKI, J. (2010), "The Dynamic Pivot Mechanism", *Econometrica*, **78**, 771–789.
- BERTSEKAS, D. (1987), *Dynamic Programming: Deterministic and Stochastic Models*. Englewood Cliffs, NJ: Prentice Hall.
- 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**, 992–1026.
- BREIMAN, L. (1968), *Probability*. Reading, MA: Addison-Wesley.
- ÇELEN, B. and KARIV, S. (2004), "Distinguishing Informational Cascades from Herd Behavior in the Laboratory", *American Economic Review*, **94**, 484–498.
- CHAMLEY, C. and GALE, D. (1994), "Information Revelation and Strategic Delay in a Model of Investment", *Econometrica*, **62**, 1065–1085.

- CHAMLEY, C. P. (2004), *Rational Herds: Economic Models of Social Learning*. Cambridge: Cambridge University Press.
- DOW, J. (1991), "Search Decisions with Limited Memory", *Review of Economic Studies*, **58**, 1–14.
- DOYLE, M. (2010), "Informational Externalities, Strategic Delay, and Optimal Investment Subsidies", *Canadian Journal of Economics*, **43**, 941–966.
- EASLEY, D. and KIEFER, N. (1988), "Controlling a Stochastic Process with Unknown Parameters", *Econometrica*, **56**, 1045–1064.
- EYSTER, E. and RABIN, M. (2010), "Naïve Herding in Rich-information Settings", *American Economic Journal: Microeconomics*, **4**, 221–243.
- EYSTER, E. and RABIN, M. (2014), "Extensive Imitation is Irrational and Harmful", *Quarterly Journal of Economics*, **129**, 1861–1898.
- GAGNON-BARTSCH, T. and RABIN, M. (2016), "Naive Social Learning, Mislearning, and Unlearning" (Harvard mimeo).
- GITTINS, J. C. (1979), "Bandit Processes and Dynamical Allocation Indices", *Journal of the Royal Statistical Society, Series B*, **41**, 148–177.
- GLAZER, J., KREMER, I. and PERRY, M. (2015), "Crowd Learning without Herding: A Mechanism Design Approach" (Warwick mimeo).
- HERRERA, H. and HÖRNER, J. (2012), "A Necessary and Sufficient Condition for Information Cascades" (Columbia and Yale mimeo).
- KAMENICA, E. and GENTZKOW, M. (2011), "Bayesian Persuasion", *American Economic Review*, **101**, 2590–2615.
- KARLIN, S. (1968), *Total Positivity*, Vol. 1. Stanford, CA: Stanford University Press.
- KARLIN, S. and RUBIN, H. (1956), "The Theory of Decision Procedures for Distributions with Monotone Likelihood Ratio", *The Annals of Mathematical Statistics*, **27**, 272–299.
- MARCH, C. and ZIEGELMEYER, A. (2020), "Altruistic Observational Learning", *Journal of Economic Theory*, **190**, 105123.
- MARSHALL, A. W. and OLKIN, I. (1979), *Inequalities: Theory of Majorization and Its Applications*. San Diego: Academic Press.
- MEDRANO, L. A. and VIVES, X. (2001), "Strategic Behavior and Price Discovery", *Rand Journal of Economics*, **32**, 221–248.
- PASTORINO, E. and KEHOE, P. J. (2011), "Individual and Market Experimentation with Multi-armed Dependent Bandits" (FRB of Minneapolis mimeo).
- QUAH, J. K.-H. and STRULOVICI, B. (2009), "Comparative Statics, Informativeness, and the Interval Dominance Order", *Econometrica*, **77**, 1949–1992.
- RADNER, R. (1962), "Team Decision Problems", *Annals of Mathematical Statistics*, **33**, 857–881.
- ROTHSCHILD, M. (1974), "A Two-armed Bandit Theory of Market Pricing", *Journal of Economic Theory*, **9**, 185–202.
- SMITH, L. and SØRENSEN, P. (2000), "Pathological Outcomes of Observational Learning", *Econometrica*, **68**, 371–398.
- TIAN, J. (2014), "Banks Inequality and Monotone Comparative Statics" (Wisconsin mimeo).
- VIVES, X. (1997), "Learning from Others: A Welfare Analysis", *Games and Economic Behavior*, **20**, 177–200.