

Informational Herding, Optimal Experimentation, and Contrarianism^{*†}

Lones Smith[‡]
Department of Economics
University of Wisconsin

Peter Norman Sørensen[§]
Department of Economics
University of Copenhagen

Jianrong Tian[¶]
School of Economics and Finance
University of Hong Kong

(under revision for the *Review of Economic Studies*)

July 20, 2017

Abstract

In a herding model, people ignore future informational gains from their action choices. We characterize the constrained efficient outcome by solving a social planner’s problem as a Bayesian optimal experimentation problem. We find that (a) herding is socially efficient, but should occur less readily, since cascade sets shrink; (b) efficiency entails *contrarian* behaviour — i.e. individuals should optimally lean against taking the myopically more popular actions, under a robust new information log-concavity condition; (c) the efficient outcome can be decentralized, rewarding individuals if their successor mimics their action, and, (d) our sufficient condition precludes finite-time cascades among purely self-interested individuals.

^{*}This supersedes “Informational Herding and Optimal Experimentation” by the first two authors, based on Chapter 2 of Sørensen’s 1996 thesis.

[†]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 the MIT theory lunch, the Stockholm School of Economics, the Stony Brook Summer Festival on Game Theory (1996), Copenhagen University, the 1997 European Winter Meeting of the Econometric Society (Lisbon), Bristol University, Erasmus University, and the 2005 Workshop on Informational Herding (Copenhagen) for comments on various versions. Smith gratefully acknowledges financial support for this work from NSF grants SBR-9422988 and SBR-9711885, and Sørensen equally thanks the Danish Social Sciences Research Council.

[‡]e-mail address: lones@ssc.wisc.edu

[§]e-mail address: peter.sorensen@econ.ku.dk

[¶]e-mail address: jt2016@hku.hk

Contents

1	INTRODUCTION	1
2	THE FORWARD-LOOKING HERDING MODEL	6
3	DYNAMIC PROGRAMMING AND CONVEX DUALITY	7
4	SHRINKING CASCADE SETS VIA PATIENCE	9
5	COMMUNICATION VIA ACTION CHOICES	12
6	CONTRARIANISM	14
6.1	An Illustrative Example: The Professor and his Student	14
6.2	Monotone Posterior Beliefs: Cascades Cannot Start Late	15
6.3	Contrarian Behavior	17
6.4	The Detailed Proof of Contrarianism with Two Actions	19
7	IMPLEMENTATION	20
A	VALUE FUNCTIONS: PROOF OF LEMMA 1	22
B	INCREASING PATIENCE PROOFS	23
B.1	Cascade Sets: Proof of Lemma 2 and More	23
B.2	Strict Value Monotonicity: Proof of Lemma 3	24
B.3	Cascade Sets and Impatience: Proof of Proposition 2(a)	25
B.4	Increased Patience and Herding: Proof of Proposition 2(b)	26
B.5	Strict Convexity of Value Function: Proof of Lemma 5	28
B.6	Natural Action Order: Proof of Corollary 2	29
C	CONTRARIANISM PROOFS	29
C.1	Bellman Derivative Formula: Proof of Lemma 6	29
C.2	Subtangents to a Convex Function: Proof of Lemma 7	30
C.3	Contrarianism: Proof of Proposition 3 for Multiple Actions	30
C.4	Strict Contrarianism: Proof of Corollary 4	32
D	IMPLEMENTATION PROOFS	32
D.1	Implementation: Proof of Proposition 4	32
D.2	Mimicry with Two Actions: Proof of Corollary 5	33

E	THREE EXAMPLES	33
E.1	Calculations for the Binomial Signal Example in §4	33
E.2	An Example of Actions in Non-Natural Order	34
E.3	The Role of Posterior Monotonicity in Contrarianism	35

1 INTRODUCTION

Conformity can arise when people collectively learn by seeing each other’s actions. This was first formalized in herding models by Bikhchandani, Hirshleifer, and Welch (1992) and Banerjee (1992). They let a sequence of privately informed individuals with identical preferences choose actions from a common finite menu, after seeing all prior actions in order. Eventually, a herd arises — over time, people settle on the same choice.

This is an important economic setting that should ring familiar. It involves an informational externality, since actions partially convey hidden private signals, but myopic individuals do not account for the value of signaling information to successors. For the observable actions of predecessors offers a trail of breadcrumbs for others to follow. It is then natural to wonder about the welfare properties. This was Banerjee’s primary issue: “1. The equilibrium pattern of choices may be (and for a large enough population, *will be*) inefficient in the ex ante welfare sense.”¹ In this paper, we definitively address this issue. We formulate the social planner’s optimization of the discounted sum of expected utilities of all agents, formally addressing the inefficiency of the equilibrium. To keep our analysis manageable, we consider the workhorse model with two states and a finite action set.

In his proposed remedy for the externality, Banerjee suggested excluding early agents from viewing others’ actions, thereby rendering their actions independent signals; Sgrou (2002) more carefully explored this idea. We find that the social planner can do better. The planner should act like an experimenter who dictates how posterior beliefs should map into actions. He trades off the myopic payoffs and the signaling value of actions. Both considerations lead him to maintain the character of the herding model, whereby an agent takes an action for an interval of posterior beliefs. Our main findings show that socially optimal behavior demands that everyone err on the side of choosing more informative actions *in every period*. This marginally discourages mimicry, but does not preclude it.

In Smith and Sørensen (2000), we showed that public beliefs (the planner’s beliefs, here) almost surely converge to a *cascade set*, where actions reflect no private information. In this case, a herd happens, and all signals lead to the same action. If private signals have uniformly bounded precision, then the cascade sets of the two extreme actions are non-trivial intervals. Our first main finding establishes that this remains true in the planner’s solution: If selfish individuals succumb to errant herds, then it is inefficient to entirely

¹Banerjee (1992) imagined a continuum of states and actions. Our analysis posits finite spaces like Bikhchandani, Hirshleifer, and Welch (1992), and so does not strictly apply to this version of the model.

preclude them. Proposition 2 shows that the cascade sets strictly shrink in the discount factor, and at some point, the non-extreme cascade sets vanish altogether; in the perfect patience limit, the two extreme sets converge to 0 and 1. The patient planner who values the welfare of later individuals still allows herds, but simply waits for more extreme beliefs. In other words, we prove that herds and inefficient herds owe not to the selfishness of agents, but to the problem of signaling private information through finitely many actions.

In our second major substantive finding, we explore how the social planner optimally skews behavior for any fixed discount factor. In Proposition 3, we argue that under general conditions, people should be encouraged to act in a *contrarian* fashion, leaning against trending popular actions and relying more on private signals. More precisely, the threshold posterior belief separating two actions — constant for myopic agents — should increase in the public belief. This penalizes extreme action for extreme public beliefs.

We show how to implement the social planner’s optimum using a new Vickrey-Clarke-Groves mechanism. As is well-known, such mechanisms internalize the externalities of action choices. Even though this externality obviously depends on the unobserved state of the world, Proposition 4 shows how the transfer to an agent need only be conditioned on both his and his successor’s actions. This scheme works because the successor’s action is informative of the true state. With just two actions, the incentive scheme rewards anyone who is mimicked by their successor (Corollary 5). Notably, even though the planner wishes to discourage mimicry, individuals are optimally rewarded by others mimicking them. In the same spirit as academic citations, the social planner efficiently punishes conformity with *predecessors’* actions (contrarian logic) by rewarding emulation by *successors*.

These are the substantive conclusions of our model. But our paper makes a several technical contributions that are generally useful in social learning and experimentation.

1. At the heart of our paper is a social planner’s optimal experimentation exercise. His state variable is the public belief based on the action history. The optimization balances current payoffs and future benefits from better information. This is not a bandit problem; rather, it is a recommender system to someone playing a standard finite action choice. Indeed, unlike the multiarmed bandit solved by the Gittins’ index, the action payoffs here are not independent — for all payoffs are impacted by the same unobserved state of the world. Nevertheless, we develop a tractable extension of this index logic.

Proposition 1 rewrites the Bellman equation optimization as a choice among finitely many *welfare indexes*. Each index is a linear function of the current agent’s posterior belief, namely, a subtangent to the value function. Each subtangent translates the value

of that action to any predecessor as a function of his own posterior belief. An individual selects the action with the highest welfare index for his posterior. With welfare indexes linear in posterior beliefs, the planner’s problem is analogous to the myopic problem — and each action is optimal for an interval of posterior beliefs. But the planner’s optimum may invoke myopically dominated actions, or may map private signals to actions in a myopically suboptimal fashion (see §5). Here we see the importance of the pure signaling value of an action, since the optimal communication pattern might not be monotone.

2. Bayesian updating is subtle in social learning precisely because one learns from actions and not signals. For one might think that conditional on seeing a fixed high action, the posterior public belief should increase in the prior public belief. But this intuition fails, because the observation of the high action no longer offers as strong an endorsement of the high state. For instance, with a public belief 0.48, the high action might be optimal when the private signal exceeds 0.52, but with a 0.52 public belief, the private signal need only exceed 0.48. If there is large mass of private signals in $(0.48, 0.52)$, then the final posterior can jump down. Yet our contrarianism result depends on securing the intuitive monotonicity, and so we must preclude masses in the private signals. We prove that a sufficient condition for a monotone relation between the prior and posterior public belief is that the unconditional distribution of the private log-likelihood ratios has a log-concave density (Lemma 4). This is the general condition we alluded to earlier for contrarianism.

This result is also important for all social learning papers. For in light of Smith and Sørensen (2000), *no cascade starts after the first period in the standard herding model given this robust signal property* (Corollary 3). A vast array of papers exploiting the possibility of sparking a cascade crucially rely on a non-monotone map from public to posterior beliefs. This clarifies that the assumption of the multinomial distribution in Bikhchandani, Hirshleifer, and Welch (1992) was not without loss of generality, but instead was one of the very few standard distributions for which cascades could eventually start.

3. Pursuing any comparative static in dynamic programming, such as the contrarian result, is immensely complicated. For lacking a twice differentiable value function, we cannot apply the Implicit Function Theorem. Barring this standard approach, we are aware of no general sensitivity analysis for that does not exploits some single crossing condition — which also does not apply. We instead exploit convex duality to deduce a single-crossing property about subtangents to a convex function. So equipped by this and the monotone map from prior to posterior public beliefs, we deduce how the planner’s welfare indexes shift. This allows us to exploit a recent advance in monotone comparative statics methods

by Quah and Strulovici (2009) to prove our contrarianism result, Proposition 3.

RELATED LITERATURE. We think that Banerjee was the first to propose a remedy for the social learning externality. To reiterate, our planner has Banerjee’s remedy at his disposal, but it is dynamically suboptimal.² Centrally planned social learning is a topical and important problem, and new optimal mechanisms have recently been explored in applied settings by Glazer, Kremer, and Perry (2015) and Che and Hörner (2016).³

The planner’s optimum is also a *team equilibrium* (Radner, 1962), where everyone seeks to maximize the sum of discounted expected utilities. There do however exist less efficient equilibria among these altruistic agents, because successors cannot fully interpret a deviation by an agent who chooses an unanticipated map of private signals into actions.⁴

Vives (1997) studies a team social learning problem in the market setting. He proves that team members choose to reveal more of their private information.⁵ Our more elaborate contrarianism comparative statics result finds that teams shy away *more* from the *more* popular actions.⁶ Vives also finds that the optimal long-run Gaussian precision growth is as low as in the selfish model. This may seem analogous to our finding that cascade sets have a non-empty interior in the team setting, but he never finds incomplete learning.

In Dow (1991), 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.

The incorrect herding outcome is intuitively related to the familiar failure of complete learning in optimal experimentation. Rothschild’s (1974) analysis of the two-armed bandit

²Closer to our spirit, Doyle (2010) considers the social planner’s problem in the endogenous-timing herding model of Chamley and Gale (1994).

³Among other papers, in Glazer, Kremer, and Perry (2015), the action includes an unobserved decision to acquire information. They also provide a good discussion of the tangential relation of our models to those in Ely (2017), Ely, Frankel, and Kamenica (2015), Hörner and Skrzypacz (2016), Kamenica and Gentzkow (2011), Kremer, Mansour, and Perry (2014) and Rayo and Segal (2010).

⁴The experimental evidence of March and Ziegelmeyer (2016) confirms contrarian behaviour among altruistic agents. Ali and Kartik (2012) explore herding in a sequential action model with a different form of payoff externalities. Anderlini, Gerardi, and Lagunoff (2012) assume that team members can communicate costlessly, but do not have perfectly aligned interests.

⁵In 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.

⁶Vives always employs the normal learning model, ruling out results like ours on the distributional shape’s importance. On the other hand, that model allows the long-run properties of learning to be characterized by the speed with which the precision approaches infinity. Our analysis offers no analogy.

is a classic example: An impatient monopolist optimally experiments with two possible prices each period, with fixed uncertain purchase chances for each price. Rothschild showed that (i) the monopolist eventually settles down on one price, and (ii) with positive probability, that eventual price is not the most profitable. This is analogous to (i) an action herd occurs, and (ii) with positive chance is ex-post incorrect. Yet, the analogy is subtle. Easley and Kiefer (1988) prove that complete learning generically arises in experimentation problems with finite state and action spaces. This is puzzling, since the herding outcome arises in a model with finite actions and states.

The formulation of our social planner’s problem offers a resolution of this puzzle. Even though each agent chooses from a finite action set, our social planner has no access to private signals, and so cannot dictate the choice among any two actions. Rather, for each history, he chooses a continuously defined rule that maps agents’ signals into actions. In the myopic planner case with a zero discount factor, we obtain the original herding model. Hence, we can conclude that the herding outcome is formally equivalent to incomplete learning in an experimentation model with a continuous choice space.⁷

We have found an original economic reason for the efficiency of contrarian behaviour. To be sure, contrarianism is a widely used term in economics and finance. For instance, investment advisors often recommend contrarian investment as a privately optimal best response to the overreaction of stock markets to information (e.g. Jegadeesh (1990)). Likewise, prudential macroeconomic policy recommends leaning against excesses in financial markets (Dewatripont and Tirole, 1994). But in our model, there is no such underlying fundamental drift: Rather, public beliefs follow a martingale, and at any point in time, they provide the best estimate of the state given the observed action history. Instead, contrarianism is socially efficient for us as it facilitates revelation of private information.⁸

The paper is organized as follows. We formulate the constrained efficient discounted herding model in §2, and give our convex duality characterization of optimal behavior in §3. We show how cascade sets shrink as patience rises in §4, and §5 gives examples where the planner’s desire to generate public information strongly overrides natural features of myopic behaviour. We motivate and explore contrarianism in §6, and §7 gives the implementation results linked to our novel welfare indexes. Many proofs are appendicized.

⁷Bose, Orosel, Ottaviani, and Vesterlund (2006) and (2008) study a monopoly seller who affects the decisions of arriving buyers through the adjustment of price. Gill and SgROI (2008) and (2012) suppose the seller can provide information before buyers arrive. The seller does not maximize buyer welfare.

⁸Pastorino 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.

2 THE FORWARD-LOOKING HERDING MODEL

We start with the standard herding model of Smith and Sørensen (2000) (SS). An infinite sequence $n = 1, 2, \dots$ of *decision-makers* (*agents*) act in that exogenous order, and share a common 50-50 prior belief over two payoff relevant *states* $\omega \in \{L, H\}$.⁹

The n th agent sees a random *private signal* σ_n about the realized state. As is common, we reduce notation by identifying his signal with his resulting interim belief $\sigma_n = \Pr(H|\sigma_n)$. The signals are i.i.d. across agents in each state $\omega = L, H$, with cdf F^ω , and have common support, say $\text{supp}(F)$. No signal perfectly reveals the state, so that, as SS then argue, F^H and F^L are mutually absolutely continuous, with derivative $dF^H/dF^L = \sigma/(1 - \sigma)$. Accordingly, they are ranked $F^H(\sigma) \leq F^L(\sigma)$, with inequality strictly inside $\text{supp}(F)$.

Agents share a common utility function u over actions and states. The action a , taken from the finite action set $\{1, \dots, A\}$, yields payoff $u(a, \omega)$ in state $\omega \in \{H, L\}$. We assume that action 1 is best in state L , and action A in state H , or $u(1, L) > u(a, L)$ for all $a \neq 1$, and $u(A, H) > u(a, H)$ for all $a \neq A$. Moreover, payoffs obey increasing differences: $u(1, H) - u(1, L) < u(2, H) - u(2, L) < \dots < u(A, H) - u(A, L)$. And finally no two actions yield the same payoffs in either state, so that for no actions $a \neq a'$ do payoffs coincide, $u(a, \omega) = u(a', \omega)$ for some state ω . The chosen action a_n provides to agent n the *expected payoff*

$$\bar{u}(a, \rho) = (1 - \rho)u(a, L) + \rho u(a, H). \quad (1)$$

We allow dominated actions a , with $\bar{u}(a, \rho) < \sup_{a'} \bar{u}(a', \rho)$ for all beliefs $\rho \in [0, 1]$ on H .

Before choosing his action a_n , the n 'th agent observes σ_n and the history of $n - 1$ predecessors' actions. He can compute the probability distribution over histories, knowing his predecessors' strategies, and end at the public belief $\pi_n = \Pr(H|a_1, \dots, a_{n-1})$. At public belief π , the private signal σ has distribution $F^\pi = \pi F^H + (1 - \pi)F^L$. Combining a conditionally independent private signal σ with public belief π gives the *posterior belief*

$$\rho = r(\pi, \sigma) \equiv \frac{\pi \sigma}{\pi \sigma + (1 - \pi)(1 - \sigma)}. \quad (2)$$

This paper explores welfare properties of the herding model, maximizing the discounted sum of future payoffs. We assume that an informationally constrained social planner observes the history of actions but not private signals. A *rule* ξ for this planner dictates

⁹As in SS, a fair public prior belief simplifies exposition, and is without loss of generality — analysis of the dynamic problem covers the continuation starting from any public belief.

an action a for each private signal $\sigma \in [0, 1]$. Let Ξ be the set of all such rules.¹⁰ A *strategy* s_n for the n 'th agent assigns a rule to each action history of length $n - 1$. The planner chooses the strategy profile $s = (s_1, s_2, \dots)$ to maximize the expected average present value of utility stream u_n :

$$\sup_s E[(1 - \delta) \sum_{n=1}^{\infty} \delta^{n-1} u_n]. \quad (3)$$

The original herding model assumes $\delta=0$. We study the patient case $\delta \in (0, 1)$.

3 DYNAMIC PROGRAMMING AND CONVEX DUALITY

The social optimum can be achieved by a Markovian dynamic optimization, whose state variable is the public belief π . A stationary *policy* assigns a rule ξ to every public belief π . With this policy, starting at belief π , the continuation value of (3) is a function $v_\delta(\pi)$.

The rule ξ prescribes any action a for signals $\sigma \in \xi^{-1}(a)$, and thus with probability $\psi(a, \omega, \xi) = \int_{\xi^{-1}(a)} dF^\omega$ in state ω , and unconditional probability $\psi(a, \pi, \xi) = \int_{\xi^{-1}(a)} dF^\pi$.¹¹ This leads to the *posterior public belief* $p(a, \pi, \xi)$. Given the planner's policy s , these action probabilities and continuations fully describe the dynamic evolution of public beliefs.

When action a is taken with positive probability (namely, $\psi(a, \pi, \xi) > 0$), we call the action and its continuation belief $p(a, \pi, \xi)$ *active*, and Bayes updating then yields continuation belief $p(a, \pi, \xi) = \pi \psi(a, H, \xi) / \psi(a, \pi, \xi)$. The martingale property of beliefs implies $p(a, \pi, \xi) = \int_{\xi^{-1}(a)} r(\pi, \sigma) dF^\pi$, namely, the expected posterior (2). When $\psi(a, \pi, \xi) = 0$, Bayes rule does not identify the belief $p(a, \pi, \xi)$. In this case, we impose a weak refinement that $p(a, \pi, \xi) = r(\pi, \sigma)$ for some possible signal $\sigma \in \text{supp}(F)$.

By dynamic programming, the planner's value function v_δ solves the Bellman equation:

$$v(\pi) = \sup_{\xi \in \Xi} (T_\xi v)(\pi), \quad (4)$$

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

$$(T_\xi v)(\pi) = \sum_{a=1}^A \psi(a, \pi, \xi) [(1 - \delta) \bar{u}(a, p(a, \pi, \xi)) + \delta v(p(a, \pi, \xi))]. \quad (5)$$

¹⁰With the possibility of "atoms", or private signals σ with positive probability under F , we should also allow for mixed rules that map private signals σ into a probability distribution over actions. To simplify exposition, our text refers only to pure rules. But all results remain valid for mixed rules.

¹¹If say rule ξ mandates action 1 for private signals $\sigma \in [0, 1/2]$, then $\psi(1, \omega, \xi) = \int_0^{1/2} dF^\omega = F^\omega(1/2)$ and $\psi(a, \pi, \xi) = \pi F^H(1/2) + (1 - \pi) F^L(1/2)$. The mixed rule extension is straightforward.

The value function v_δ solving optimization (4) is convex in beliefs, since the upper envelope of linear functions is convex, as is standard in expected utility maximization. Intuitively, it equals the payoffs of state-optimal actions 1 and A at beliefs 0 and 1, and therefore its extreme slopes cannot exceed $u(1, H) - u(1, L)$ and $u(A, H) - u(A, L)$ — see Figure 1 for intuition. Finally, since the planner’s information grows over time in this otherwise stationary problem, a more patient planner enjoys a higher discounted payoff.

Lemma 1. *The value function v_δ is a bounded and continuous convex function of public beliefs π , with subtangent¹² slopes at least $u(1, H) - u(1, L)$ and at most $u(A, H) - u(A, L)$. The value is weakly increasing in the discount factor δ .*

From a dual perspective, as is well-known, the convex function v is the upper envelope of its supporting subtangent lines. We denote this *subtangent space* as $\mathcal{T}_v \subset \mathbb{R}^2$, which is compact, since subtangents are parameterized by their slope and intercepts. 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 the sum and maximization to rewrite operator (5) as

$$(T_\xi v)(\pi) = \max_{(\tau_1, \dots, \tau_A) \in \mathcal{T}_v^A} \sum_{a=1}^A \int_{\xi^{-1}(a)} [(1 - \delta)\bar{u}(a, r(\pi, \sigma)) + \delta\tau_a(r(\pi, \sigma))] dF^\pi. \quad (6)$$

Exchange the order of the sup in (3) with the max in (6) to get the planner’s dual problem:¹³

$$v(\pi) = \max_{(\tau_1, \dots, \tau_A) \in \mathcal{T}_v^A} \sup_{\xi \in \Xi} \sum_{a=1}^A \int_{\xi^{-1}(a)} [(1 - \delta)\bar{u}(a, r(\pi, \sigma)) + \delta\tau_a(r(\pi, \sigma))] dF^\pi. \quad (7)$$

In the multi-armed bandit (§6.5 of Bertsekas, 1987), an experimenter each period chooses among a finite action set, with random *independent* rewards. 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: At public belief π and private posterior ρ , the agent chooses the action a with the largest *welfare index* w — equal to the social payoff as privately gauged by the agent.

¹²When v is differentiable, this is a tangent; otherwise, this is a supporting tangent line.

¹³As an aside, convex duality offers a computational strategy for solving the dynamic programming problem. In the iterative process, given a value v_n , the next value v_{n+1} is obtained in principle by searching across all the possible rules. But the 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 .

Proposition 1. *For any public belief π , an agent with posterior ρ takes the action a with maximal welfare index*

$$w(a, \pi, \rho) = (1 - \delta)\bar{u}(a, \rho) + \delta\tau_a(\rho), \quad (8)$$

where τ_a is a supporting tangent to v at continuation belief $p(a, \pi, \xi)$.

Proof. First, (7) implies the index expression (8): the supremum over rules ξ is attained by allocating private signal σ to action a where $(1 - \delta)\bar{u}(a, r(\pi, \sigma)) + \delta\tau_a(r(\pi, \sigma))$ is maximal, and $\rho = r(\pi, \sigma)$. For a fixed rule ξ , (6) implies that τ_a is subtangent to v at $p(a, \pi, \xi)$. \square

The subtangent to the value function $\tau_a(\rho)$ embeds some important economics, allowing the planner to evaluate his payoff for any agent's hypothetical realized private posterior ρ (Figure 2). In particular, whether a higher posterior helps or hurts the payoffs of later individuals depends solely on the slope of the value function at the public belief. For the privately informed agent with belief ρ assigns value $\tau_a(\rho)$ to the continuation game after action a where followers act optimally on the public belief $p(a, \pi, \xi)$.

An *interval rule* partitions private posterior beliefs in $[0, 1]$ into possibly empty intervals $\mathcal{I} = \{I_a\}$, with action a optimal iff $\rho \in I_a$. The affine welfare indices coincide at the boundary between neighboring intervals: $w(a, \pi, \rho) = w(\bar{a}, \pi, \rho)$ on the boundary ρ between I_a and $I_{\bar{a}}$. These boundary posterior beliefs, or *thresholds*, play a critical role in our paper.¹⁴

Corollary 1. *For each public belief π , an optimal interval rule \mathcal{I} exists.*¹⁵

Proof. From (8), each action's welfare index is affine in the posterior belief. Action a thus has maximal index value on a convex set, i.e., an interval. \square

When it helps, we can also re-interpret this as an interval rule in signal space too, since the posterior belief $\rho = r(\pi, \sigma)$ continuously increases in the private signal σ .

4 SHRINKING CASCADE SETS VIA PATIENCE

As in the herding model, we focus on public beliefs where active learning stops. The public belief π lies inside the *cascade set* $C_a(\delta)$ — namely, public beliefs where a is optimal for all private signals σ . Thus, it is optimal to choose action a with probability one ($\psi(a, \pi, \xi) = 1$)

¹⁴We allow for the possibility that the rule uses mixing at the threshold.

¹⁵This result corresponds to the interval partition deduction Proposition 1 in Dow (1991), albeit with a Bayesian binary state structure not present in Dow. Meanwhile, when $\delta = 1$, our Proposition 1 more roughly corresponds to the FOC of Dow's Proposition 2.

exactly when $v(\pi) = \bar{u}(a, \pi)$; this is uniquely optimal strictly inside $C_a(\delta)$ (see Figure 1). As the myopic payoff frontier $\bar{u}(a, \pi)$ is affine in π , and the value v is convex, this equality holds on a closed interval $C_a(\delta)$. When the discount factor δ increases, v weakly increases by Lemma 1, and therefore $C_a(\delta)$ weakly shrinks. Since $\bar{u}(a, \cdot)$ is a tangent to the value v on $C_a(\delta)$, we have $w(a, \pi, \rho) = \bar{u}(a, \rho)$. The union $C(\delta) = \cup_{a=1}^A C_a(\delta)$ is the *cascade set*.

Private signals are called *unbounded* if the (compact) signal support $\text{supp}(F)$ contains 0 and 1, and are called *bounded* if $\text{supp}(F) \subseteq (0, 1)$. Smith and Sørensen (2000) show that interval cascade sets exist iff private signals are bounded. We next argue that this useful characterization result also describes the social optimum of our herding model.

Lemma 2. (a) For discount factors $\delta \in [0, 1)$, $0 \in C_1(\delta)$ and $1 \in C_A(\delta)$, and $C(\delta) \neq [0, 1]$.
(b) With bounded signals, $C_1(\delta) = [0, \underline{\pi}(\delta)]$, $C_A(\delta) = [\bar{\pi}(\delta), 1]$, 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 $a \neq 1, A$.

So cascade sets weakly shrink in the discount factor δ . But with bounded private signals, the two extreme cascade sets for actions 1 and A are nonempty for all $\delta < 1$: For near extreme public beliefs 0 and 1, Bayesian updating requires too many rounds of myopic sacrifice to change the myopically best action, and active experimentation is suboptimal.

As in Smith and Sørensen (2000), we prove in §B.1-B that beliefs converge almost surely to the cascade sets, and that learning is incomplete iff private signals are bounded. We now argue that, with bounded signals, cascade sets strictly shrink in δ (see Figure 1). As $\delta \uparrow 1$, the limit cascade sets are $\{0, 1\}$, and the chance of an incorrect herd vanishes.

Proposition 2. Assume bounded private signals.

(a) Non-empty cascade sets strictly shrink when $\delta < 1$ rises: For all actions a , if $\delta_2 > \delta_1$ and $C_a(\delta_1) \neq \emptyset$, then $C_a(\delta_2) \subset C_a(\delta_1)$. For large enough $\delta < 1$, all cascade sets disappear except for $C_1(\delta)$ and $C_A(\delta)$, while $\lim_{\delta \rightarrow 1} C_1(\delta) = \{0\}$ and $\lim_{\delta \rightarrow 1} C_A(\delta) = \{1\}$.

(b) A herd almost surely starts, and the chance it is incorrect vanishes as $\delta \uparrow 1$.

This result formalizes many economists' gut feeling that cascades are inefficient — that the more weight that is placed on future individuals, the smaller is the cascade set. The proof exploits the planner's indifference towards experimentation at the edge of a cascade set. He strictly prefers to experiment if he grows slightly more patient, and so the cascade set shrinks. Intuitively, the more patient planner enjoys a higher value of information (expected gain in the continuation value). The proof requires that we strengthen Lemma 1.

Lemma 3. The value function increases strictly in δ on $[0, 1)$ outside the cascade sets: If $\delta_2 > \delta_1$, then $v_{\delta_2}(\pi) > v_{\delta_1}(\pi)$ for all public beliefs $\pi \notin C(\delta_2)$.




Figure 1: **Myopic Payoffs, the Bellman Value, and Cascade Sets.** By Proposition 2, each cascade set $C_\delta(\pi)$ shrinks as the discount factor δ rises. Inside the cascade set, the planner's value coincides with the myopic value.

Since cascade sets shrink in the discount factor, we can explain why an equilibrium need not be a social optimum (recall the discussion in §7). For there exists a suboptimal team equilibrium when $\delta > 0$ in which everyone acts myopically, since herding on action a is a team equilibrium for public beliefs $\pi \in C_a(0) \setminus C_a(\delta)$. If every successor considers agent n 's action uninformative, as π is in a cascade set for $\delta = 0$, the best that agent n can do is to maximize his current payoff. But since $\pi \notin C_a(\delta)$, this rule is suboptimal.

We illustrate our main result here with an example similar in spirit to the original inspirational example in Bikhchandani, Hirshleifer, and Welch (1992). Assume two actions with symmetric payoffs $u(2, H) = u(1, L) = 1$, $u(1, H) = u(2, L) = -1$, and a symmetric binary private signal on the two-point support $\{\sigma_0, \sigma_1\}$, with $F^H(\sigma_0) = 1 - F^L(\sigma_0) = \sigma_0 = 1 - \sigma_1 < 1/2$. In the myopic case $\delta = 0$, for public beliefs $\pi \in (\sigma_0, \sigma_1)$, the optimal action depends on the private signal: take action 2 iff the signal outcome is σ_1 . But for extreme public beliefs, the private signal should be ignored, and one should take action 2 if $\pi > \sigma_1$ and action 1 if $\pi < \sigma_0$. We argue that the cascade set is smaller with $\delta > 0$.

By symmetry, the cascade sets are $(0, \bar{\pi}(\delta))$ and $(1 - \bar{\pi}(\delta), 1)$. The optimal strategy is simple. Active experimentation occurs for public beliefs $\pi \in (1 - \bar{\pi}(\delta), \bar{\pi}(\delta))$. Starting at any such π , we need only count the number of high signal realizations σ_1 necessary to surpass the public belief $\bar{\pi}(\delta)$, or the number of low realizations σ_0 needed to depress the public belief below $1 - \bar{\pi}(\delta)$. Then $(1 - \bar{\pi}(\delta), \bar{\pi}(\delta))$ can then be partitioned into sub-intervals, overlapping at end-points only, where this pair of numbers is constant.

Fix $\pi_3 \equiv \bar{\pi}(\delta)$. Since indifference between actions prevails at $\pi = 1/2$, and π_3 is in a

cascade, the same action is best at both signals. So π_3 exceeds one Bayes updating step¹⁶ above $1/2$. For simplicity, assume a small enough discount factor $\delta > 0$ that π_3 and $1 - \pi_3$ differ by at most three steps. Choose π_1 one step below π_3 , and π_2 two steps above $1 - \pi_3$. Then $\pi_2 > \pi_1 > 1/2$. By symmetry, $1 - \pi_2$ is two steps below π_3 and one step below π_1 , and $1 - \pi_1$ is one step above $1 - \pi_3$ and one step below π_2 .

Altogether, $[0, 1]$ is divided into seven adjacent subintervals $[0, 1 - \pi_3]$, $[1 - \pi_3, 1 - \pi_2]$, $[1 - \pi_2, 1 - \pi_1]$, $[1 - \pi_1, \pi_1]$, $[\pi_1, \pi_2]$, $[\pi_2, \pi_3]$, $[\pi_3, 1]$. Within an interval, current and future reactions to signal strings are fixed; the public belief π only describes the chance of either state, and therefore the value $v(\pi)$ is linear in π over each sub-interval. In Appendix E.1, we derive and illustrate in Figure 3 the piecewise linear value function. We find in particular that:

$$\pi_3 = \bar{\pi}(\delta) = \sigma_1 + \frac{2\delta(2\sigma_1 - 1)^2\sigma_0\sigma_1}{(2 - \delta)(1 - 2\delta^2\sigma_0\sigma_1) + \delta(2\sigma_1 - 1)^2}. \quad (9)$$

*This solution is increasing and continuous in the discount factor δ , and equal to σ_1 at $\delta = 0$. Thus, the cascade set $[\pi_3, 1]$ shrinks in the discount factor δ . The analysis is valid when π_3 is less than two steps over $1/2$, and thus for small enough $\delta > 0$.*¹⁷

That $\pi_3 > \sigma_1$ in this example captures the public good aspect of experimentation. For indeed, when the public belief lies in $[\sigma_1, \pi_3)$, agents take the myopically dominated low action for a low signal outcome, leaning against their myopic interests.

5 COMMUNICATION VIA ACTION CHOICES

In the socially planned herding model, the communication value of actions can overwhelm myopic payoff considerations. Inspired by the search model of Dow (1991), we make two observations about the planner’s actions choices reflecting this insight.¹⁸

LESSON 1: DOMINATED ACTIONS MAY BE SOCIALLY VALUABLE.

Learning is filtered through a finite mesh action screen in a herding model. Welfare would intuitively rise if agents could more finely convey their private information via more actions. More available actions creates a larger alphabet to communicate. This suggests that using slightly myopically dominated actions might be efficient.

¹⁶Since the signal is symmetric and binary, we can use the word “steps” to be upward or downward Bayesian updates via (2) that cancel when applied together.

¹⁷Our method extends to larger $\delta < 1$. But the simplicity is deceptive. With an asymmetric signal, Bayesian updating does not produce equal up and down steps, and the value function is not piecewise affine.

¹⁸We 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.

To see this, suppose that action A dominates $A - 1$, with $u(A, \omega) = u(A - 1, \omega) + \varepsilon$.¹⁹ Suppose that private signals are bounded, $\text{supp}(F) \subset (0, 1)$. Then, by Lemma 2 (b), there is an interval cascade set $C_A = [\bar{\pi}, 1]$, where $\bar{\pi} < 1$. We claim that action $A - 1$ is optimal with positive probability for some public beliefs for small $\varepsilon > 0$ and/or a large enough discount factor $\delta < 1$. For if not, the planner never uses action $A - 1$. Since $v(\pi) = \bar{u}(A, \pi)$ iff $\pi \in C_A$, the value function v is not locally linear near $\bar{\pi}$. At the public belief $\bar{\pi}$, action A is optimal for all private signals. Now, consider an alternative rule that maps smaller private signals $\sigma \leq 1/2$ into action $A - 1$, and larger private signals $\sigma > 1/2$ into action A . This induces continuation public beliefs $p(A - 1, \bar{\pi}, \xi) < \bar{\pi} < p(A, \bar{\pi}, \xi)$. Since the convex value function is not linear near $\bar{\pi}$, 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 the discount factor δ is so large that $\delta\eta > (1 - \delta)\varepsilon$.

LESSON 2: THE ACTION ORDER MIGHT NOT BE MYOPICALLY OPTIMAL.

The *natural order* requires that the interval $I_{a'}(\delta)$ lie above $I_a(\delta)$ if the actions $a' > a$ are both active. By our increasing differences assumption, a myopic agent uses this order. But for a high enough discount factor $\delta < 1$, short-run payoff considerations do not dominate, and the optimal intervals can differ somewhat from the myopically optimal ones. Once this happens, the optimal map between actions and intervals need no longer be natural, as we illustrate by example in Appendix E.2. But such non-natural orders require a high enough discount factor δ . To underscore the essential role of patient agents, let us bound the discount factor away from one. By supermodularity, the payoff slope differences $\Delta_a \equiv (u(a, H) - u(a, L)) - (u(a - 1, H) - u(a - 1, L)) > 0$ for actions $a = 2, \dots, A$. Define the sum $\Delta \equiv (u(A, H) - u(A, L)) - (u(1, H) - u(1, L))$ and minimum $\underline{\Delta} = \min_{\{2, \dots, A\}} \Delta_a$.

Corollary 2. *If $\delta < \underline{\Delta}/(\Delta + \underline{\Delta})$, then for any public belief π not in a cascade set, the optimal rule uses the natural action ordering. With two actions, this holds for $\delta < 1/2$.*

The impatience premise of Corollary 2 is essential. In the example in Appendix E.2, the unnatural action order is optimal for all large enough discount factors $\delta < 1$.

¹⁹ Strictly speaking, these actions do not have increasing differences. The illustration is simpler this way, but it would also suffice to let $u(A - 1, L) = u(A, L) - \varepsilon/2$.

6 CONTRARIANISM

6.1 An Illustrative Example: The Professor and his Student

We illustrate contrarian behavior and its optimality in a fully-solved two period example. This is a pure information transmission problem like Dow (1991), with $\delta = 1$ (full altruism).

A professor and his student share a common prior π on state H , and observe private signals σ , with cdf's $F^H(\sigma) = \sigma^2$ and $F^L(\sigma) = 2\sigma - \sigma^2$. The professor sees the signal, and either takes action 1 or 2; his student observes the professor's action, but not his signal. Subject to this restriction, the professor selflessly acts to maximize his student's expected payoff, where his state payoffs are $u(2, H) = u(1, L) = 1, u(1, H) = u(2, L) = -1$.

If the student starts with a continuation public belief p , then she takes action 2 exactly when her signal $\sigma \geq 1 - p$. Now, $\sigma \geq 1 - p$ with chance $1 - F^H(p) = 1 - (1 - p)^2$ in state H and with chance $F^L(1 - p) = p^2$ in state L . Hence, the student's value function is

$$V_S(p) = p(1 - (1 - p)^2 - (1 - p)^2) + (1 - p)(1 - p^2 - p^2) = 1 - 2p + 2p^2.$$

The professor uses an interval partition rule, equivalent to choosing action 1 for low signals $\sigma < \bar{\sigma}$, and action 2 if $\sigma \geq \bar{\sigma}$.²⁰ He seeks to maximize $V(\pi) = E[V_S(P)|\pi]$, where P is his student's random public belief, and the expectation is taken before the private signal σ is realized. Since $\pi = E[P|\pi]$ by the martingale property of beliefs, we have

$$V(\pi) = E[V_S(P)|\pi] = E(1 - 2P + 2P^2|\pi) = 1 - 2\pi + 2\pi^2 + 2E[(P - \pi)^2|\pi].$$

Then the professor's optimal value $V(\pi)$ exceeds the student value $V_S(\pi) = 1 - 2\pi + 2\pi^2$ by twice the variance of beliefs. This variance encapsulates the option value of learning in this two period setting. We now compute this term. Given the threshold rule, a different continuation public belief $P = p_1$ or $P = p_2$ arises after each of the two professorial actions 2 and 1. Bayes rule reveals the formulas $p_1(\bar{\sigma}) = [\pi\bar{\sigma}^2]/[\pi\bar{\sigma}^2 + (1 - \pi)(2\bar{\sigma} - \bar{\sigma}^2)]$ and $p_2(\bar{\sigma}) = [\pi(1 - \bar{\sigma}^2)]/[\pi(1 - \bar{\sigma}^2) + (1 - \pi)(1 - 2\bar{\sigma} + \bar{\sigma}^2)]$. We can explicitly compute:

$$E[(P - \pi)^2|\pi] = \frac{\pi - p_1}{p_2 - p_1}(p_2 - \pi)^2 + \frac{p_2 - \pi}{p_2 - p_1}(\pi - p_1)^2 = (p_2 - \pi)(\pi - p_1). \quad (10)$$

Only this term in $V(\pi)$ depends on $\bar{\sigma}$. Maximizing (10) over $\bar{\sigma}$ yields *private signal threshold*

²⁰With no weight on current payoffs ($\delta = 1$), swapping actions 1 and 2 is also optimal. WLOG, we assume the natural ordering.

$\bar{\sigma}(\pi) = (\pi - 1 + \sqrt{\pi - \pi^2}) / (2\pi - 1)$ if $\pi \neq 1/2$, with limit $\bar{\sigma}(1/2) = 1/2$ by l'Hopital's rule.

Corresponding to this private signal threshold, the professor's optimal *posterior belief threshold* is

$$\theta(\pi) = r(\pi, \bar{\sigma}(\pi)) = \frac{\pi - \sqrt{\pi - \pi^2}}{2\pi - 1}$$

Illustrating our later short-run contrarianism result, $\theta(\pi)$ increases in π . So the professor optimally communicates the state of the world by acting in a “contrarian” fashion. He leans against the public belief, choosing action 2 less often when state H is more probable.²¹

6.2 Monotone Posterior Beliefs: Cascades Cannot Start Late

Before deriving contrarianism in our infinite horizon problem with $\delta < 1$, we first address a basic monotonicity property of Bayesian updating in the finite-action model. We derive a robust condition on the private signal distribution yielding *posterior monotonicity*: holding fixed an interval partition \mathcal{I} for posterior beliefs, the continuation public belief²² $p(a, \pi, \mathcal{I})$ strictly increases in the current public belief π , for all active actions a .

Given the equi-likely states, the *unconditional signal distribution* is described by the function $F = (F^H + F^L)/2$. When the density $f = F'$ exists, Bayesian updating implies a simple “no introspection condition” (Smith and Sørensen, 2000) for densities: $f^H(\sigma) = 2\sigma f(\sigma)$ and $f^L(\sigma) = 2(1 - \sigma)f(\sigma)$. Associate to private signal σ the log-likelihood ratio $\ell = \Lambda(\sigma) \equiv \log(\sigma/(1 - \sigma))$, with inverse $\mathcal{S}(\ell) = e^\ell / (1 + e^\ell)$. *In the rest of the paper, we maintain the following (novel) regularity condition.*^{23,24}

(LC): *The log-likelihood ratio density $\phi(\ell) \equiv f(\mathcal{S}(\ell))\mathcal{S}'(\ell)$ exists, and is log-concave.*

Assumption (LC) is violated by atomic distributions, but common continuous distributions are log-concave (see Marshall and Olkin (1979), §18.B.2.d), including that in §6.1.

That the Bayes rule map from π and σ into ρ is linear in log-likelihood ratios, namely $\Lambda(\rho) = \Lambda(\pi) + \Lambda(\sigma)$, underlies our next result. Let $g(\rho|\pi)$ be the posterior belief density.

Lemma 4. *Posterior monotonicity obtains for any active action, because the posterior belief density $g(\rho|\pi)$ is strictly log-supermodular, given (LC).*

²¹It is easy to verify that the posterior threshold $\theta(\pi)$ exceeds the constant, myopic posterior belief threshold $\theta(\pi) = 1/2$ exactly when $\pi > 1/2$.

²²Abusing notation, we let $p(a, \pi, \mathcal{I})$ denote the continuation belief when the rule \mathcal{I} is fixed as π varies.

²³Smith and Tian (2016) develop more general conditions for posterior monotonicity.

²⁴Recent work by Morris and Yildiz (2016) derives a related nonmonotone map from priors to posteriors, for which they venture a sufficient fat tails assumption (that precludes log-concavity).

Proof: Action a is taken for posteriors $\rho \in I_a$, and with chance $\int_{I_a} g(\rho|\pi)d\rho > 0$, since a is active. The continuation public belief is thus $p(a, \pi, \mathcal{I}) = \int_{I_a} \rho g(\rho|\pi)d\rho / \int_{I_a} g(\rho|\pi)d\rho$. We next show that g obeys the MLRP (log-supermodular), and so $p(a, \pi, \mathcal{I})$ increases in π .

Let $\phi^\omega(\ell)$ be the density over the private log likelihood ratio ℓ in state ω . Observe that $\phi(\ell) = (\phi^L(\ell) + \phi^H(\ell))/2 = (1 + e^\ell)\phi^L(\ell)/2$, as the no introspection condition implies $\phi^H(\ell) = e^\ell\phi^L(\ell)$. Hence, $\log \phi^L(\ell) = \log \phi(\ell) - \log(1 + e^\ell) + \log 2$ is strictly concave, since ϕ is log-concave by condition (LC). As a result, the unconditional density $h(\ell|\pi)$ over the posterior log likelihood ratios ℓ given prior belief π is strictly log-supermodular, given:

$$h(\ell|\pi) = (1 - \pi)\phi^L(\ell - \Lambda(\pi)) + \pi\phi^H(\ell - \Lambda(\pi)) = (1 - \pi)(1 + e^\ell)\phi^L(\ell - \Lambda(\pi))$$

Then $g(\rho|\pi)$ is strictly log-supermodular, as the map $\ell \mapsto \rho$ strictly increases. \square

To see the role played in Lemma 4 by the log-concavity assumption (LC), we now give an example in which posterior monotonicity fails when (LC) fails. We slightly modify our signal family example in §6.1, punching a hole in its support. Choose $b \in (1/2, 1)$, and define the density $f(\sigma) = 1/(2 - 2b)$, for $\sigma \leq 1 - b$ and $\sigma \geq b$, and $f(\sigma) = 0$ otherwise. Let $f^H(\sigma) = 2\sigma f(\sigma)$ and $f^L(\sigma) = 2(1 - \sigma)f(\sigma)$. Suppose action 2 is optimal if the posterior belief exceeds $1/2$. Given a public prior belief $\pi > b$, the posterior likelihood ratio after seeing action 2 is

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

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

When $\delta = 0$, every action a is taken for posterior beliefs in a fixed interval $[\theta_1, \theta_2]$. In their “bounded beliefs example”, Smith and Sørensen (2000) found that public beliefs can transition into a cascade set if and only if the posterior public belief after an action is not monotone in the prior belief. That is, landing in a cascade set requires that (LC) fail:

Corollary 3. *Given assumption (LC), a cascade cannot start after period one if $\delta = 0$.*

It is instructive to observe that the multinomial signal examples with cascades in the seminal paper by Bikhchandani, Hirshleifer, and Welch (1992) violate assumption (LC).²⁵

²⁵Herrera and Hörner (2012) note for the binary action model that posterior monotonicity is equivalent to two properties being true: the increasing hazard ratio property and the increasing failure ratio property. They copy arguments from Smith and Sørensen (2000) to note that posterior monotonicity precludes cascades. They incorrectly claim that the property is necessary.

6.3 Contrarian Behavior

We now consider short run contrarian behavior, that individuals increasingly lean against actions increasingly favoured by popular beliefs. We generalize the logic of the professor-student example, and show that under the log-concavity assumption (LC), the posterior belief threshold separating pairs of actions satisfies $\theta(\pi) < \theta(\pi')$ for any non-cascade public belief realizations $\pi < \pi'$. The generalization allows more than two actions, an infinite-horizon with discounted future payoffs, and the possibility of multiple optimal rules.

We have observed 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 both π and π' , and that the set of active actions is identical.

Fixing one such action order, re-label the active actions so that higher actions are taken at higher signals (we simply call the number of such actions A). Let θ_a denote the threshold between private posterior beliefs leading to actions a and $a + 1$, and define the *threshold vector* $\theta = (\theta_1, \dots, \theta_{A-1})$. The *threshold space* $\Theta(\pi) \subset \mathbb{R}^{A-1}$ is the set of vectors θ where $r(\min \text{supp}(F), \pi) < \theta_1 < \dots < \theta_{A-1} < r(\max \text{supp}(F), \pi)$. For an interval rule defined by the vector θ , the probability of action a is $\psi(a, \pi, \theta)$, and the continuation belief $p(a, \pi, \theta)$.

Let $\Theta^*(\pi) \subset \Theta(\pi)$ be the set of optimal threshold vectors. We formally call behaviour *contrarian* if, for any pair $\pi < \pi'$ with this identical optimal action ordering, the set $\Theta^*(\pi')$ is higher than $\Theta^*(\pi)$ in the strong set order.²⁶ This coincides with an intuitive notion of first order stochastic dominance: at the higher belief π' , any set of lower actions $\{1, \dots, a\}$ is taken for a higher set of posterior beliefs $[0, \theta'_a]$. Behaviour is *strictly contrarian* if, for all $\theta \in \Theta^*(\pi)$ and $\theta' \in \Theta^*(\pi')$, we have $\theta' \gg \theta$, so that all coordinates of θ' are higher than θ .

Proposition 3. *Given (LC), behaviour is contrarian for any discount factor $\delta \in [0, 1)$.*

Figure 2 depicts an instructive key step of the proof of Proposition 3. The threshold between actions a and $a + 1$ equates the welfare indexes $w(a, \pi, \theta_a) = w(a + 1, \pi, \theta_a)$; also, $w(a, \pi, \rho)$ down-crosses $w(a + 1, \pi, \rho)$ at the posterior belief $\rho = \theta_a$. In other words, the net gain to taking the higher action grows in the posterior belief ρ .

Next, by the formula (8), the public belief π affects the welfare index w only through the value function tangent $\tau_a(\rho)$. And at the higher public belief $\pi' > \pi$, the tangents shift along the convex value function as shown in Figure 2, forcing $w(a, \pi', \theta_a) > w(a + 1, \pi', \theta_a)$.²⁷ The posterior threshold where $w(a, \pi, \theta_a) = w(a + 1, \pi, \theta_a)$ then rises, as desired: $\theta'_a > \theta_a$.

²⁶Recall that Y' dominates Y in the strong set order if $y \in Y$ and $y' \in Y' \Rightarrow y \vee y' \in Y'$ and $y \wedge y' \in Y$.

²⁷The proof works for any convex continuation value function — it does not require our infinite horizon dynamic optimization. So Proposition 3 applies to the two-period professor-student problem in §6.1.




Figure 2: **How Tangents Comove.** The tangent to the convex value function at any public belief π measures the present value to all later individuals starting at any posterior belief (Proposition 1), and thus higher beliefs raise this value iff the value function slopes up. Given the continuation posterior threshold θ_a , we draw the tangents at posteriors $p_a(\pi) < p_{a+1}(\pi)$. By Lemma 7, as π rises, so do $p_a(\pi)$ and $p_{a+1}(\pi)$, while the tangent $\tau_{a+1}(\theta_a)$ falls and the tangent $\tau_a(\theta_a)$ rises.

Assumption (LC) guarantees updating monotonicity of public beliefs (Lemma 4), and thereby the noted monotone tangent difference, by the above logic. Contrarianism can fail without monotone public beliefs, as we show in Appendix E.3.

We now strengthen Proposition 3, and secure strictly contrarian behavior. By the above logic, this holds if $\delta > 0$ and the value function is strictly convex. For then the tangents in (8) have positive weight, and they strictly shift in π .

Corollary 4. *If $\delta \in (0, 1)$, signals obey assumption (LC), and all actions are taken in the natural order, then behavior is strictly contrarian outside the cascade sets.*

The result holds if the value function is strictly convex in a neighborhood of any continuation belief. But it is affine on an interval $[\underline{z}, \bar{z}]$ whenever there exists a constant optimal strategy on $[\underline{z}, \bar{z}]$ — in particular, the value function is affine on cascade sets.²⁸ We next claim that the cascade sets are in fact the only affine portions of the value function.²⁹

Lemma 5. *If the signal support $\text{supp}(F)$ is convex and all actions are taken in the natural order, then the value function v is strictly convex outside the cascade sets.*

²⁸A strategy, started at $\pi \in [\underline{z}, \bar{z}]$, yields some state-contingent expected values v^H and v^L . The expected value of following the same strategy, starting at belief ρ , is then $\tau(\rho) = (1 - \rho)v^L + \rho v^H$. Since the strategy is optimal at π and feasible at ρ , the affine τ is tangent to v at π . If v is affine, then $v(\rho) = \tau(\rho)$, and the strategy is optimal for all $\rho \in [\underline{z}, \bar{z}]$. Conversely, if the strategy is optimal, $v(\rho) = \tau(\rho)$ for all $\rho \in [\underline{z}, \bar{z}]$.

²⁹Note that Assumption (LC) implies a convex belief support.

6.4 The Detailed Proof of Contrarianism with Two Actions

We now explain the local argument of Proposition 3 with $A = 2$, because it is instructive.

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) with (5) 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 (5):

$$B(\theta|\pi) = \sum_{a=1}^2 \psi(a, \pi, \theta) [(1 - \delta)\bar{u}(a, p(a, \pi, \theta)) + \delta v(p(a, \pi, \theta))]. \quad (11)$$

Solutions to the constrained problem $\max_{\theta \in \Theta(\pi)} B(\theta|\pi)$ define an optimizer set $\Theta^*(\pi)$. To prove Proposition 3, it suffices that $\Theta^*(\pi)$ increase in the strong set order.

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 (8) and (11), we derive an expression for $B_\theta(\theta|\pi)$ in terms of the welfare index.

Lemma 6. *The Bellman function B is differentiable almost everywhere with derivative*

$$B_\theta(\theta|\pi) = g(\theta, \pi) (w(1, \pi, \theta) - w(2, \pi, \theta)). \quad (12)$$

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

The next result is a useful property of tangents to a convex function (refer to Figure 2).³⁰

Lemma 7. *Fix $z_1 < z_2 < z_3$ and a convex function v . Let τ_i be a tangent function to the value function v at z_i . Then $\tau_2(z_1) \geq \tau_3(z_1)$ (respectively, $\tau_1(z_3) \leq \tau_2(z_3)$), with strict inequality unless v is affine on $[z_2, z_3]$ (respectively, on $[z_1, z_2]$).*

Returning to our proof of Proposition 3, suppose that the thresholds $\theta \in \Theta^*(\pi)$ and $\theta' \in \Theta^*(\pi')$ are inversely ordered as $\theta' < \theta$ — otherwise, we're done. Since $r(\sigma, \pi)$ is an increasing function of 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 (12) weakly

³⁰Recalling Proposition 1, τ_i would correspond to the value function at continuation belief z_i .

increases in the public belief π , when $\tilde{\theta} \in [\theta', \theta]$. By Lemma 4, continuation beliefs rise in public beliefs: $p(a, \pi', \tilde{\theta}) > p(a, \pi, \tilde{\theta})$ for $a = 1, 2$. The two cases in Lemma 7 yield, as desired,

$$\Delta(\tilde{\theta}, \pi') - \Delta(\tilde{\theta}, \pi) = \delta\{[\tau'_1(\tilde{\theta}) - \tau_1(\tilde{\theta})] + [\tau_2(\tilde{\theta}) - \tau'_2(\tilde{\theta})]\} \geq 0. \quad (13)$$

Next, $\alpha(\tilde{\theta}) \equiv g(\tilde{\theta}|\pi')/g(\tilde{\theta}|\pi)$ is a positive and nondecreasing function over $[\theta', \theta]$, since g is log-supermodular, by Lemma 4. Then Lemma 6 and inequality (13) imply:

$$B_\theta(\tilde{\theta}|\pi') = g(\tilde{\theta}|\pi')\Delta(\tilde{\theta}, \pi') \geq g(\tilde{\theta}|\pi')\Delta(\tilde{\theta}, \pi) = \alpha(\tilde{\theta})B_\theta(\tilde{\theta}|\pi), \quad (14)$$

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.

Now consider the stronger claim in Corollary 4 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 our 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 $w(2, \pi', \rho) - w(1, \pi', \rho)$ increases in ρ , as follows from (8). For the natural action order implies that $\bar{u}(2, \rho) - \bar{u}(1, \rho)$ is strictly increasing, and convexity of v implies that its tangent difference $\tau'_2(\rho) - \tau'_1(\rho)$ 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 5. Lemma 7 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 Lemma 5. \square

7 IMPLEMENTATION

Can the planner implement the optimal solution using a feasible transfer scheme for the selfish agents? Since he cannot observe the private signals, transfers would have to depend on the observed action history alone. Otherwise, given the index formula (8), the socially optimal behavior can be decentralized by awarding individuals transfers $\delta\tau_a(\rho)/(1 - \delta)$ depending on the posterior belief ρ .

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 ordering of actions, and thus are not sufficient if the selfish optimal action ordering is not socially optimal.³¹ We solve this using richer transfers.

A *pivot mechanism* that rewards agents for their marginal contribution to social welfare, i.e. from changing the public belief, would align the agents’ and planner’s incentives. So we would need to pay agents the incremental discounted value $M(a, \pi, \omega)$ of successors from observing the current action a in each state ω .³² Now, even though the planner does not know the state ω , future agents’ choices rely on *their* information, and so indirectly on ω . We use this linkage to implement the social outcome with transfers $t(a, b)$ that depend just on the current and next agent’s actions a and b . Thus, an agent with posterior belief ρ expects to receive a premium $\rho M(a, \pi, H) + (1 - \rho)M(a, \pi, L)$ from taking action a .

Proposition 4. *The social optimum (3) can be implemented for selfish agents by a mechanism whose transfers only depend on the public belief, and actions of the agent and his successor. A unique such mechanism exists if no continuation belief is in a cascade set.*

We next argue that with the myopic action order and two actions, the pivot mechanism transfers in Proposition 4 reward individuals who are mimicked by successors. Intuitively, an action is more likely smart if the successor’s signal leads him to emulate it.

Corollary 5. *Assume the myopic action ordering in the binary action world. The transfers are ranked $t(a, a) \geq t(a, b)$ whenever $b \neq a$ and neither belief π nor $p(a, \pi, \xi)$ are in $C(\delta)$.*

Related to the implementation results for selfish agents, consider the problem where everyone altruistically aims to maximize a welfare measure. Adapting Radner (1962), we call a perfect Bayesian equilibrium of this game a *team equilibrium*. 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 some agent n has a strictly better reply $\tilde{\xi}$ at a history. Then the planner can improve his value at that history by *fully* mimicking this deviation, i.e. (i) using rule $\tilde{\xi}$ in the first period and then (ii) 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 profitable deviation contradicts optimality of s .

This decentralization result squares well with our observation after Proposition 1 that the welfare index w is the altruistic player’s value at posterior belief ρ . Note however, that other team equilibria also exist, as we will explain after Lemma 3 in §4.

³¹We could restrict the planner to such rules that cannot implement our planner’s solution. Bru and Vives (2002) likewise consider IC mechanisms that cannot implement the optimum of Vives (1997).

³²The formal definition of $M(a, \pi, \omega)$ is given in §D.1 with the proof of Proposition 4.

A VALUE FUNCTIONS: PROOF OF LEMMA 1

We use the Bellman operator $T = \sup_{\xi \in \Xi} T_\xi$ from the RHS of (4). From (4) and (5), 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, convex functions. We simply 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:

$$p(a, \pi_\lambda, \xi) = \frac{\lambda\psi(a, \pi_1, \xi)}{\psi(a, \pi_\lambda, \xi)}p(a, \pi_1, \xi) + \frac{(1-\lambda)\psi(a, \pi_2, \xi)}{\psi(a, \pi_\lambda, \xi)}p(a, \pi_2, \xi). \quad (15)$$

The first (myopic) term in (5) 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 (15) holds, the second (future) term obeys:

$$\psi(a, \pi_\lambda, \xi)v(p(a, \pi_\lambda, \xi)) \leq \lambda\psi(a, \pi_1, \xi)v(p(a, \pi_1, \xi)) + (1-\lambda)\psi(a, \pi_2, \xi)v(p(a, \pi_2, \xi)). \quad (16)$$

Then $Tv(\pi_\lambda) = T_\xi v(\pi_\lambda) \leq \lambda T_\xi v(\pi_1) + (1-\lambda)T_\xi v(\pi_2) \leq \lambda Tv(\pi_1) + (1-\lambda)Tv(\pi_2)$, by summing (16) over actions $a = 1, \dots, A$.

Let $\tilde{u}(\pi) = \max_a \bar{u}(a, \pi)$ denote the payoff frontier. The bound on tangent slopes follows from the observations that $v(0) = u(1, L)$ and $v(1) = u(A, H)$, that the convex function v exceeds the payoff frontier \tilde{u} , and that $\bar{u}(1, \rho)$ and $\bar{u}(A, \rho)$ define the most extreme slopes of \tilde{u} , by supermodularity.

Claim 1. *The function sequence $\{T^n \tilde{u}\}$ pointwise increases and converges to v_δ . The value v_δ weakly exceeds \tilde{u} , and strictly so outside the cascade sets.*

Proof. To maximize $\sum_{a=1}^A \psi(a, \pi, \xi) [(1-\delta)\bar{u}(a, p(a, \pi, \xi)) + \delta\tilde{u}(p(a, \pi, \xi))]$ over rules $\xi \in \Xi$ for the given belief π , one rule $\tilde{\xi}$ a.s. chooses the myopically optimal action. Then $p(\tilde{\xi}(\sigma), \pi, \tilde{\xi}) = \pi$ a.s., resulting in value $\tilde{u}(\pi)$. Optimizing over all $\xi \in \Xi$, we get $T\tilde{u}(\pi) \geq \tilde{u}(\pi)$ for all π . By induction, $T^n \tilde{u} \geq T^{n-1} \tilde{u}$, yielding a pointwise increasing sequence converging to the fixed point $v_\delta \geq \tilde{u}$. Finally, when π is outside the cascade sets, by definition it is *not* optimal to induce one action a.s., whence $v_\delta(\pi) > \tilde{u}(\pi)$ if $\delta \in [0, 1)$. \square

Claim 2. *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 $v \geq \tilde{u}$. If δ increases, then $T_\xi \tilde{u}$ pointwise increases too, since more weight

is placed on the larger component of the RHS of (5). By (4), $T\tilde{u}$ is pointwise higher. Iterating this argument, $T^n\tilde{u}$ is higher. Let $n \rightarrow \infty$ and apply Claim 1. \square

B INCREASING PATIENCE PROOFS

B.1 Cascade Sets: Proof of Lemma 2 and More

A. PROOF OF LEMMA 2. For (a), note that action 1 is myopically strictly optimal when $\pi = 0$. Since it updates to continuation belief $\pi = 0$ for any rule, it is also dynamically optimal for any discount factor $\delta \in [0, 1)$. A similar proof holds for $\pi = 1$. Since the private signal is valuable in the selfish problem, $\cup_{a=1}^A C_a(0) \neq [0, 1]$.

(b) For low public beliefs, it is optimal to let the rule ξ induce 1; the argument for high beliefs is similar. Action 1 is optimal at belief $\pi = 0$, and there is no tie, so 1 is the optimal selfish choice for beliefs $\pi \leq \pi'$, 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 beliefs π in the interval $[0, \pi'/2]$. No action can reveal a stronger private signal than any $\sigma \in \text{supp}(F) \subseteq [\sigma_1, \sigma_0] \subset (0, 1)$. 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 of v_δ , $v_\delta(\bar{p}(\pi)) - v_\delta(\pi)$ is less than $\eta(1 - \delta)/\delta$ for small enough π . By the Bellman equation (4), any action $a \neq 1$ is strictly suboptimal for such small beliefs.

(c) Assume unbounded signals. Smith and Sørensen (2000) prove that $C_a(0) = \emptyset$ for all $a \neq 0, 1$, and that $C_1(0) = \{0\}$ and $C_A(0) = \{1\}$. \square

B. CASCADE SETS AS LIMIT BELIEFS. As in Smith and Sørensen (2000), public beliefs converge by the martingale convergence theorem, and the limit is not fully wrong:

Claim 3. *The belief process $\langle \pi_n \rangle$ is a martingale unconditional on the state, converging a.s. to some limiting random variable π_∞ . The limit π_∞ is concentrated on $(0, 1]$ in state H .*

Smith and Sørensen (2000) find for $\delta = 0$ that the public belief process converges upon the cascade set. The result extends to the case $\delta > 0$:

Theorem 1. *Consider a solution of the planner's problem. The limit belief π_∞ has support in $C_1(\delta) \cup \dots \cup C_A(\delta)$. In particular, π_∞ is concentrated on the truth for unbounded signals.*

Proof: At least two actions occur with positive chance for any belief π not in any cascade set. By the interval structure of Corollary 1, the highest such action is more likely in

state H , and the lowest in state L . So the continuation belief differs from π with positive probability. Intuitively, or by the characterization result for Markov-martingale processes in Appendix B of Smith and Sørensen (2000), π cannot lie in the support of π_∞ .

B.2 Strict Value Monotonicity: Proof of Lemma 3

Fix $\delta_2 > \delta_1$. Fix $\pi \notin C(\delta_2)$. If $\pi \in C(\delta_1)$, we're done, since $v_{\delta_1}(\pi) = \tilde{u}(\pi) < v_{\delta_2}(\pi)$. If $\pi \notin C(\delta_1)$, the δ_1 -optimal rule ξ induces with positive chance some action \hat{a} with continuation belief $p(\hat{a}, \pi, \xi) \notin C(\delta_1)$. To see why, recall that π is the average of the continuation beliefs, and that $C_a(\delta_1)$ is an interval. Moreover, at most one cascade set is hit (Claim 5 below). 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 , with strict inequality for \hat{a} , since $\delta_2 > \delta_1$. By (4) and (5), the δ_1 -optimal rule ξ provides a strictly higher value than $v_{\delta_1}(\pi)$, for the discount factor δ_2 . Optimizing over rules for δ_2 , we conclude that $v_{\delta_2}(\pi) > v_{\delta_1}(\pi)$. \square

Claim 4. *For any $\delta > 0$ and any action $a \in \{1, \dots, A\}$, if $\tilde{\pi} \in (0, 1)$ is an endpoint of cascade set $C_a(0)$ then $\tilde{\pi} \notin C_a(\delta)$.*

Proof: Let $\tilde{\pi} = \min C_a(0)$ where $a \neq 1$. Denote the minimal posterior belief by $\check{\rho} = r(\tilde{\pi}, \min \text{supp}(F))$. Then $\bar{u}(a - 1, \check{\rho}) = \bar{u}(a, \check{\rho})$. Define

$$w_{a-1}(\rho) = (1 - \delta)\bar{u}(a - 1, \rho) + \delta\tau(\rho) \quad \text{and} \quad w_a(\rho) = (1 - \delta)\bar{u}(a, \rho) + \delta\bar{u}(a, \rho) \quad (17)$$

where τ is a tangent of v_δ at $\check{\rho}$. Since $\check{\rho} < \tilde{\pi}$, we have $\check{\rho} \notin C_a(0)$. As noted before, Lemma 1 implies $C_a(\delta) \subseteq C_a(0)$, so $\check{\rho} \notin C_a(\delta)$. Thus, $\bar{u}(\check{\rho}, a) < v_\delta(\check{\rho}) = \tau(\check{\rho})$. Plugging this inequality into (17) gives $w_a(\check{\rho}) < w_{a-1}(\check{\rho})$. If $\tilde{\pi} \in C_a(\delta)$, then $w_a(\check{\rho})$ is the welfare index at posterior belief $\check{\rho}$, but our inequality then contradicts Proposition 1. \square

This yields a general new property of cascade sets:

Claim 5. *Let $\delta > 0$. For any belief π , continuation beliefs lie in at most one cascade set.*

Proof: Given unbounded signals, continuation beliefs never lie in a cascade set. Assume bounded signals. Let $\underline{\sigma} = \min \text{supp}(F)$ and $\bar{\sigma} = \max \text{supp}(F)$. Suppose that for some π , two continuation beliefs $\pi_1 < \pi_2$ lie in distinct cascade sets, namely, $C_{a'}(\delta)$ below $C_{a''}(\delta)$. Then $\pi_1 \in C_{a'}(0)$ and $\pi_2 \in C_{a''}(0)$ since, again, $C_a(\delta) \subseteq C_a(0)$ for any a . Let $\pi' = \max C_{a'}(0) \leq \pi'' = \min C_{a''}(0)$. Then $\pi_1 \leq \pi'$. There exist x_1, x_2 in $[\underline{\sigma}, \bar{\sigma}]$ with $r(\pi, x_1) = \pi_1$

and $r(\pi, x_2) = \pi_2$. Since (a) Bayes-updating commutes, (b) $r(\pi, \bar{\sigma}) \geq r(\pi, x_2) = \pi_2$ and $x_1 \geq \underline{\sigma}$, (c) $\pi_2 \geq \pi''$, and (d) $\pi'' \in C_{a''}(0)$ while $\pi' \in C_{a'}(0)$:

$$r(\pi_1, \bar{\sigma}) = r(r(\pi, x_1), \bar{\sigma}) = r(r(\pi, \bar{\sigma}), x_1) \geq r(\pi_2, \underline{\sigma}) \geq r(\pi'', \underline{\sigma}) \geq r(\pi', \bar{\sigma})$$

and so $\pi_1 \geq \pi'$. Thus $\pi_1 = \pi'$, which contradicts Claim 4. \square

B.3 Cascade Sets and Impatience: Proof of Proposition 2(a)

We now show strict inclusion of the cascade sets. Fix any action \hat{a} . Since $C_{\hat{a}}(\delta) = \{\pi | v_{\delta}(\pi) - \bar{u}(\hat{a}, \pi) = 0\}$ is closed by continuity, we prove that if $\hat{\pi} \equiv \min C_{\hat{a}}(\delta_1)$ (i.e. the left end of a cascade set) then $\hat{\pi} \notin C_{\hat{a}}(\delta_2)$ when $\hat{a} \neq 1$. (The case $\max C_{\hat{a}}(\delta_1)$ is similar.)

CASE 1: MULTIPLE OPTIMIZERS. Assume that at $\hat{\pi}$ and δ_1 , some optimal rule involves actions other than \hat{a} with positive chance. Then some continuation beliefs fall outside $C(\delta_1)$, with positive chance. For otherwise, by Claim 5, all continuation beliefs lie in the same cascade set — and this rule incurs a myopic cost (not playing \hat{a}) with no informational gain. As in the proof of Claim 4, $v_{\delta_2}(\hat{\pi}) > v_{\delta_1}(\hat{\pi}) = \bar{u}(\hat{a}, \hat{\pi})$. We conclude that $\hat{\pi} \notin C(\delta_2)$.

CASE 2: UNIQUE OPTIMIZER. Assume that the unique optimal rule with $\delta = \delta_1$ at belief $\hat{\pi}$ is to play \hat{a} at probability 1. Let $\pi_n \uparrow \hat{\pi}$. For each n , let $T_n = (\tau_{a_n}^n, a = 1, \dots, A)$ be optimal tangents in (7). Since $\hat{\pi} = \min C_{\hat{a}}(\delta_1)$, we must have $\pi_n \notin C_{\hat{a}}(\delta_1)$. Therefore, we can find some action $a'_n \neq \hat{a}$ and private signal $\sigma_n \in \text{supp}(F)$ with higher index: $(1 - \delta_1)\bar{u}(a'_n, r(\pi_n, \sigma_n)) + \delta_1\tau_{a'_n}^n(r(\pi_n, \sigma_n)) \geq (1 - \delta_1)\bar{u}(\hat{a}, r(\pi_n, \sigma_n)) + \delta_1\tau_{\hat{a}}^n(r(\pi_n, \sigma_n))$. Since $\mathcal{T}_v^{|A|}$ and $\text{supp}(F)$ are compact and A is finite, there is a subsequence where T_n has limit $T^* = (\tau_{a'}^*, a_n \text{ has limit } a' \neq \hat{a}, \text{ and } \sigma_n \text{ has limit } \hat{\sigma})$. Write $\hat{\rho} = r(\hat{\pi}, \hat{\sigma})$. By the Theorem of the Maximum, T^* is an optimal subtangent vector for $\hat{\pi}$, so

$$(1 - \delta_1)\bar{u}(a', \hat{\rho}) + \delta_1\tau_{a'}^*(\hat{\rho}) \geq (1 - \delta_1)\bar{u}(\hat{a}, \hat{\rho}) + \delta_1\tau_{\hat{a}}^*(\hat{\rho}). \quad (18)$$

First, since choosing \hat{a} with chance 1 is the unique optimal rule at $\hat{\pi}$ for δ_1 , $\tau_{\hat{a}}^*$ must be a subtangent line of v_{δ_1} at $\hat{\pi}$, i.e., $\tau_{\hat{a}}^*(\hat{\pi}) = v_{\delta_1}(\hat{\pi}) = \bar{u}(\hat{a}, \hat{\rho})$. Next, since $\hat{\pi} \in C_{\hat{a}}(\delta_1)$, by Claim 4, $\hat{\pi}$ is strictly inside $C_{\hat{a}}(0)$. So action \hat{a} is myopically strictly dominant, i.e. $\bar{u}(\hat{a}, \hat{\rho}) > \bar{u}(a', \hat{\rho})$. Then (18) implies that $\tau_{a'}^*(\hat{\rho}) > \tau_{\hat{a}}^*(\hat{\rho})$. Since $\delta_2 > \delta_1$, (18) further implies that $(1 - \delta_2)\bar{u}(a', \hat{\rho}) + \delta_2\tau_{a'}^*(\hat{\rho}) > (1 - \delta_2)\bar{u}(\hat{a}, \hat{\rho}) + \delta_2\tau_{\hat{a}}^*(\hat{\rho})$. Then optimizing over tangents for δ_2 , conclude that $v_{\delta_2}(\hat{\pi}) > \bar{u}(\hat{a}, \hat{\pi})$, and thus $\hat{\pi} \notin C_{\hat{a}}(\delta_2)$. For both $\tau_{a'}^*(\rho)$ and $\tau_{\hat{a}}^*(\rho)$ are weakly below v_{δ_2} by $\delta_1 < \delta_2$ and weak monotonicity of value.

THE PERFECT PATIENCE LIMIT. We finally prove that for actions $a \neq 1, A$, cascade sets disappear when δ tends to 1, and $\lim_{\delta \rightarrow 1} C_1(\delta) = \{0\}$ and $\lim_{\delta \rightarrow 1} C_A(\delta) = \{1\}$. Consider first any action $a \notin \{1, A\}$, and suppose δ is such that the cascade set $C_a(\delta)$ is non-empty. Let the rule ξ take $a - 1$ and a for signals σ in the respective intervals $I_{a-1} = [0, \theta]$ and $I_a = (\theta, 1]$, where $0 < F(\theta) < 1$. Updating with the optimistic news that $\sigma \in I_a$ leads to an upward revision of the public belief: There exists $\varepsilon > 0$ such that $p(a, \pi, \xi) - \pi \geq \varepsilon$ for all $\pi \in C_a(\delta) \subset (0, 1)$. Denoting by π'' the upper bound of $C_a(\delta)$, write $[\pi', \pi''] = [\pi'' - \varepsilon/2, \pi''] \cap C_a(\delta)$. Since the convex function v_δ strictly exceeds the affine $\bar{u}(a, \cdot)$ outside $C_a(\delta)$, and since $v_\delta(\pi) = \bar{u}(a, \pi)$ inside $C_a(\delta)$, there exists $\eta > 0$ so small that $\psi(a, \pi, \xi)v_\delta(p(a, \pi, \xi)) + \psi(a - 1, \pi, \xi)v_\delta(p(a - 1, \pi, \xi)) > v_\delta(\pi) + \eta$ for all $\pi \in [\pi', \pi'']$. We prove that the interval $[\pi', \pi'']$ is excised from $C_a(\delta')$ once $\delta' > \delta$ is sufficiently large. If this were not true, then $v_{\delta'}(\pi') = \bar{u}(a, \pi')$, and there is an expected gain of at least η in the continuation value of the the Bellman equation (4) by switching from the cascade rule to rule ξ . For δ' sufficiently large, this continuation gain dominates any first-period loss, proving sub-optimality of the cascade rule at π' , and hence in $[\pi', \pi'']$. By iterating this procedure a finite number of times, each time excising length $\varepsilon/2$ from interval $C_a(\delta)$, we see that $C_a(\delta)$ vanishes for large enough δ . If $a = 1$ or A , apply this procedure repeatedly: for all $\varepsilon > 0$, $C_a(\delta) \cap [\varepsilon, 1 - \varepsilon]$ vanishes for δ near 1. \square

B.4 Increased Patience and Herding: Proof of Proposition 2(b)

A herd obtains on action a at stage N if $n = N, N + 1, N + 2, \dots$ choose action a . While a cascade implies a herd, the converse is false. To show that herds arise, we extend the logic of Smith and Sørensen (2000), by extending the Overturning Principle. Claim 6 proves that for beliefs π near $C_a(\delta)$, actions other than a push the updated public belief far from its current value. The reason is that actions other than a yield a first order myopic loss against a second order learning gain if the continuation public belief is close to its current value. So belief convergence implies convergence of actions: a limit cascade implies a herd.

Theorem 2. *Consider any planner's solution:*

- (i) *A herd eventually starts.*
- (ii) *With unbounded signals, the herd is on the ex post optimal action.*
- (iii) *The chance of an incorrect herd with bounded signals vanishes as $\delta \uparrow 1$.*

The proof uses a claim that extends the Overturning Principle in Smith and Sørensen (2000) to this forward-looking model. When π is near $C_a(\delta)$, action a should occur with

high chance. More precisely, any other actions distinctly shift beliefs, or there was a non-negligible probability of observing some other action which would distinctly shift beliefs.

Claim 6. For $\delta \in [0, 1)$, assume $C_a(\delta) \neq \emptyset$. Then there exists $\varepsilon > 0$ and an ε -neighbourhood $K \supset C_a(\delta)$, such that $\forall \pi \in K \cap (0, 1)$, either:

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

Proof of Theorem 2: We cite the extended (conditional) Second Borel-Cantelli Lemma in Corollary 5.29 of Breiman (1968): Let Y_1, Y_2, \dots be any stochastic process, and events \mathcal{A}_n be measurable with respect to (Y_1, \dots, Y_n) . Then almost surely

$$\{\mathcal{A}_n \text{ infinitely often (i.o.)}\} = \left\{ \sum_{n=1}^{\infty} P(\mathcal{A}_{n+1} | Y_n, \dots, Y_1) = \infty \right\}.$$

Now, fix an optimal policy, and let Υ denote this map from public beliefs to rules, $\xi = \Upsilon(\pi)$. Choose $\varepsilon > 0$ to satisfy Claim 6 for all actions $1, 2, \dots, A$. For fixed a , define events $\mathcal{B}_n = \{\pi_n \text{ is } \varepsilon\text{-close to } C_a(\delta)\}$, $\mathcal{C}_n = \{\psi(a, \pi_n, \Upsilon(\pi_n)) < 1 - \varepsilon\}$, and $\mathcal{D}_{n+1} = \{|\pi_{n+1} - \pi_n| > \varepsilon\}$. If $\mathcal{B}_n \cap \mathcal{C}_n$ is true, then scenario (ii) in Claim 6 obtains, and so $P(\mathcal{D}_{n+1} | \pi_n) \geq \varepsilon/A$. Then $\sum_{n=1}^{\infty} P(\mathcal{D}_{n+1} | \pi_1, \dots, \pi_n) = \infty$ conditional on $\mathcal{B}_n \cap \mathcal{C}_n$ i.o. By the above Borel-Cantelli Lemma, a.s. \mathcal{D}_n obtains i.o. conditional on $\mathcal{B}_n \cap \mathcal{C}_n$ i.o. But since $\langle \pi_n \rangle$ a.s. converges by Claim 3, \mathcal{D}_n i.o. is a zero chance event, and thus so is $\mathcal{B}_n \cap \mathcal{C}_n$ i.o.

Consider the event \mathcal{E} that $\langle \pi_n \rangle$ has a limit in $C_a(\delta)$ and $\mathcal{B}_n \cap \mathcal{C}_n$ occurs only finitely often. By definition, \mathcal{E} implies that eventually $\mathcal{B}_n \setminus (\mathcal{C}_n \cup \mathcal{D}_{n+1})$. But $\mathcal{B}_n \setminus \mathcal{C}_n$ implies that every $a' \neq a$ leads to \mathcal{D}_{n+1} , by Claim 6 (i). Action a is then eventually taken on \mathcal{E} . Sum over all a to get a chance one event, by Claim 3, Theorem 1, and Proposition 2 (a). \square

Proof of Claim 6: Choose $\eta > 0$ small enough such that for any π sufficiently close to $C_a(\delta)$, we have $\psi(a', \pi, \Upsilon(\pi)) < 1 - \eta$ for any $a' \neq a$. If such η does not exist, since the optimal rule correspondence is u.h.c., almost surely taking action a' is optimal at some $\tilde{\pi} \in C_a(\delta)$. This is impossible, as a' incurs a strict myopic loss, and captures no information gain.

First, assume bounded signals. By (b) of Lemma 2, for π close enough to 0 or 1, the *only* optimal rule is to stop learning. Thus, we need only consider π in some closed subinterval I of $(0, 1)$. Let $\sigma_0 = \min \text{supp}(F)$ and $\sigma_1 = \max \text{supp}(F)$. By the existence of informative signals, $\sigma_0 < 1/2 < \sigma_1$. Let $\varepsilon > 0$ be the minimum of η , $F^H((2\sigma_0 + 1)/4)$, and $1 - F^L((2\sigma_1 + 1)/4)$ (notice that $(2\sigma + 1)/4$ is the midpoint between σ and $1/2$).

Assume $\psi(a, \pi, \Upsilon(\pi)) \geq 1 - \varepsilon$ for some $\pi \in I$. By Corollary 1, any action $a' \neq a$ is a.s. only taken for signals within either $[\sigma_0, (2\sigma_0 + 1)/4]$ or $[(2\sigma_1 + 1)/4, \sigma_1]$. Any such a'

implies case (i) of the Claim (selecting, if necessary, ε even smaller).

If instead $\psi(a, \pi, \Upsilon(\pi)) < 1 - \varepsilon$, then each action is taken with chance less than $1 - \varepsilon$. By construction of ε , different actions are taken at the two extreme signals (by the interval structure of the optimal rule). At least one of the A actions occurs with chance at least ε/A , does not include signals near $1/2$, and therefore moves public beliefs by at least ε (selecting, if necessary, ε even smaller), as claimed in case (ii) of the Claim.

Next consider unbounded signals. Let the absolute slope of the value function v have upper bound κ . Since no two payoffs are tied at 0, there exists a small $\zeta > 0$ such that the myopic action payoffs $\bar{u}(a, \rho)$ maintain the same ranking, and the difference $|\bar{u}(\tilde{a}, \rho) - \bar{u}(\check{a}, \rho)|$ exceeds $\kappa\zeta$ for all $\tilde{a} \neq \check{a}$, for all $\rho \in [0, \zeta]$.

Assume that π is near the cascade set $\{0\}$ — the other case, $\{1\}$, is similar. Then only one a'' can have low continuation belief $p(a'', \pi, \Upsilon(\pi)) \in [0, \zeta]$. If not, consider the altered policy that redirects signals from two such actions into the myopically higher of the two. This yields a first-period payoff gain of more than $\kappa\zeta$, and a future value loss of at most $\kappa\zeta$ (for p remains in $[0, \zeta]$). So the altered policy is a strict improvement.

Assume $\psi(1, \pi, \Upsilon(\pi)) \geq 1 - \varepsilon$. Then $p(1, \pi, \Upsilon(\pi)) \leq \pi/(1 - \varepsilon) \leq \zeta$, for small enough π and ε . As only action $a'' = 1$ has continuation belief in $[0, \zeta]$, case (i) is satisfied.

Finally, assume $\psi(1, \pi, \Upsilon(\pi)) < 1 - \varepsilon$. Then $\psi(a'', \pi, \Upsilon(\pi)) < 1 - \varepsilon$. Otherwise, $a'' \neq 1$ and a myopic gain of at least $(1 - \varepsilon)\zeta - \varepsilon U$ obtains from swapping the signals for 1 and a'' , without any change in future value (here U denotes the maximal possible myopic payoff difference). Thus there is a gain if ε is small enough: contradiction. Since $\psi(a'', \pi, \Upsilon(\pi)) < 1 - \varepsilon$ there must exist some other action taken with chance at least ε/A yielding continuation belief outside $[0, \zeta]$. Thus, case (ii) holds. \square

B.5 Strict Convexity of Value Function: Proof of Lemma 5

Let $\hat{\pi} \notin C(\delta)$, and suppose to the contrary that v is affine around $\hat{\pi}$. Let $\hat{\xi}$ be an optimal rule mapping private signals to actions. Define $H(\pi) = \sum_a \psi(a, \pi, \hat{\xi})w(a, \hat{\pi}, p(a, \pi, \hat{\xi}))$. First, $v(\hat{\pi}) = H(\hat{\pi})$, by (4) and Proposition 1. Second, H is affine because indices are affine, and by the martingale property of posteriors, (15). Third, $H(\pi) \leq v(\pi)$ from (7), since H employs both the particular rule $\hat{\xi}$ and the particular tangents to v at $p(a, \hat{\pi}, \hat{\xi})$.

Since v is affine, $H(\pi) = v(\pi)$ around $\hat{\pi}$. Again, (7) implies that $\hat{\xi}$ is optimal for π close to $\hat{\pi}$ and that $w(a, \hat{\pi}, \rho)$ are optimal index functions at π . To see the contradiction to this, consider two active actions. Since the actions are naturally ordered, we have

$p(1, \hat{\pi}, \hat{\xi}) < p(2, \hat{\pi}, \hat{\xi})$, and that $\bar{u}(2, \rho) - \bar{u}(1, \rho)$ is strictly increasing. The indices $w(1, \hat{\pi}, \rho)$ and $w(2, \hat{\pi}, \rho)$ then intersect at a unique point, θ . The signal support is convex and both actions are active, so posterior threshold θ uniquely characterizes the optimal rule $\hat{\xi}$ at $\hat{\pi}$. However, for $\pi \neq \hat{\pi}$ but close to $\hat{\pi}$, the fixed-private-signal-threshold rule $\hat{\xi}$ selects different actions for ρ near θ depending on π , because (2) strictly increases in π . \square

B.6 Natural Action Order: Proof of Corollary 2

By Proposition 1, it is optimal to choose the action with highest welfare index $w(a, \pi, \rho)$. Since $w(a, \pi, \rho)$ is linear in ρ , it suffices that $\partial w(a, \pi, \rho)/(\partial \rho)$ strictly increase in a . With v convex, the slope of any subtangent line τ of v is sandwiched as follows:

$$u(1, H) - u(1, L) \leq v'(0+) \leq \frac{\partial \tau}{\partial \rho} \leq v'(1-) \leq u(A, H) - u(A, L).$$

This inequality allows us to bound the difference of welfare indices (8) from below:

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

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

C CONTRARIANISM PROOFS

C.1 Bellman Derivative Formula: Proof of Lemma 6

From (11), 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 \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 use Proposition 1 to rewrite (5) as follows, proving Lemma 6:

$$B(\theta|\pi) = \int_0^\theta w(1, \pi, \rho)g(\rho|\pi)d\rho + \int_\theta^1 w(2, \pi, \rho)g(\rho|\pi)d\rho. \quad (19)$$

Claim 7. *Let $\theta \in \Theta(\pi)$. Assume $\theta_a = \dots = \theta_{a+j} = x$ for some $a \geq 1$ and $j \geq 0$ with $a+j \leq A-1$, and suppose that $\theta_{a-1} < x < \theta_{a+j+1}$.³³ Then the Bellman function B in (21)*

³³We use the notation $\theta_0 = r(\min \text{supp}(F), \pi)$ and $\theta_A = r(\max \text{supp}(F), \pi)$.

is absolutely continuous with respect to x , and its derivative in x almost everywhere equals:

$$B_x(\theta|\pi) \equiv g(x|\pi) (w(a, \pi, x) - w(a + j + 1, \pi, x)). \quad (20)$$

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 omitted proof of this many action generalization follows closely on Lemma 6, since we take action a for $\rho \in [\theta_{a-1}, x]$, and action $a + j + 1$ for $\rho \in [x, \theta_{a+j+1}]$. So the derivative of the Bellman function B in x is similar to (12) which had payoffs and tangents for actions $a = 1$ and $a + j + 1 = 2$. Thus, (20) follows. The inequality follows similarly from (14). \square

C.2 Subtangents to a Convex Function: Proof of Lemma 7

When v is affine on $[z_1, z_2]$, subtangents τ_1 and τ_2 can coincide, with $\tau_1(z_3) = \tau_2(z_3)$. Otherwise, the subtangent τ_2 is steeper than τ_1 . Thus, $\tau_2(z_3) - \tau_2(z_2) > \tau_1(z_3) - \tau_1(z_2)$, whence $\tau_2(z_3) - \tau_1(z_3) > \tau_2(z_2) - \tau_1(z_2)$. Since v is convex, the subtangent τ_1 lies below v at z_2 , so that $\tau_2(z_2) = v(z_2) \geq \tau_1(z_2)$. Hence, $\tau_2(z_3) > \tau_1(z_3)$. The z_1 analysis is similar. \square

C.3 Contrarianism: Proof of Proposition 3 for Multiple Actions

Claim 8. *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 7 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 (20), 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. \square

Fixing the action ordering, the Bellman function (5) for a convex continuation value v is:

$$B(\theta|\pi) = \sum_{a=1}^A \psi(a, \pi, \theta) [(1 - \delta)\bar{u}(a, p(a, \pi, \theta)) + \delta v(p(a, \pi, \theta))]. \quad (21)$$

We now prove Proposition 3 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 7, 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 8). Our result follows if:

$$B(\theta|\pi) - B(\theta \wedge \theta'|\pi) \geq 0 (> 0) \implies B(\theta \vee \theta'|\pi') - B(\theta'|\pi') \geq 0 (> 0). \quad (22)$$

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 (22). Conversely, if $\theta \wedge \theta'$ is not optimal at π , then θ' is not optimal at π' , by the strict inequality in (22).

We split the proof of (22) 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:

$$B(\hat{\theta}|\pi') - B(\hat{\theta} \wedge \theta'|\pi') \geq (> 0) \implies B(\theta \vee \theta'|\pi') - B(\theta'|\pi') \geq (> 0). \quad (23)$$

Then (22) follows if we also argue:

$$B(\theta|\pi) - B(\theta \wedge \theta'|\pi) \geq (> 0) \implies B(\hat{\theta}|\pi') - B(\hat{\theta} \wedge \theta'|\pi') \geq (> 0). \quad (24)$$

We now prove (24). First, for all $\theta'' \in [\hat{\theta} \wedge \theta', \hat{\theta}]$, we have $\hat{\theta} = \theta \vee \theta''$ and so:

$$B(\hat{\theta}|\pi) - B(\theta''|\pi) \geq B(\theta|\pi) - B(\theta \wedge \theta''|\pi) \geq 0, \quad (25)$$

by supermodularity of $B(\cdot|\pi)$ and optimality of θ at π , respectively. When $\theta'' = \hat{\theta} \wedge \theta'$ in (25), 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 (24). \square

C.4 Strict Contrarianism: Proof of Corollary 4

Pick $\pi' > \pi$. Let $\theta \in \Theta^*(\pi)$ and $\theta' \in \Theta^*(\pi')$. By Proposition 3, behavior is contrarian. Suppose for a contradiction that it is not strictly so, and thus $\theta'_k \leq \theta_k$ for some k . By Proposition 3, $\theta \vee \theta'$ 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 neighboring pair of active actions. \square

D IMPLEMENTATION PROOFS

D.1 Implementation: Proof of Proposition 4

For all π , let $v(\pi, \omega)$ denote the ω -contingent continuation value of the subgame starting at public belief π . (From Proposition 1, there is a tangent τ to the value function at π such that $v(\pi, L) = \tau(0)$ and $v(\pi, H) = \tau(1)$.) At given π , the additional ω -contingent present value of current action a to later agents equals

$$M(a, \pi, \omega) = [\delta/(1 - \delta)](v(p(a, \pi, \xi), \omega) - v(\pi, \omega)). \quad (26)$$

In a cascade, $\pi \in C_a(\delta)$, no transfer is needed. Since $C_a(\delta) \subseteq C_a(0)$, the selfish agent takes a . We thus continue under the assumption $\pi \notin C(\delta)$.

Suppose first that no active action reaches the cascade set. Consider active action a . The successor takes some action b for the lowest private signals, with chance denoted $\psi(b, \omega)$. Give the agent transfer $t(a, b)$ when the successor chooses b , and $t(a, -b)$ otherwise. The ω -conditional expected transfer for a is $\psi(b, \omega)t(a, b) + (1 - \psi(b, \omega))t(a, -b)$. Lower signals are more likely in the low state: $\psi(b, H) < \psi(b, L)$; therefore, there exist a unique pair $t(a, b), t(a, -b)$ solving the two equations (for $\omega = H, L$):

$$\psi(b, \omega)t(a, b) + (1 - \psi(b, \omega))t(a, -b) = M(a, \pi, \omega). \quad (27)$$

Simple algebra confirms that $\bar{u}(a, \rho) + \rho M(a, \pi, H) + (1 - \rho)M(a, \pi, L)$ is an affine transformation of the index $w(a, \pi, \rho)$, where the transformation only depends on π and δ . The transfer thus provides the right incentive and implements social optimum, and constitutes

a pivot mechanism,³⁴ since each agent is paid the marginal contribution.

We deter agents from taking inactive actions with large negative transfers.

Suppose next that the continuation belief after action a lands in a cascade set. Then $\psi(b, H) = \psi(b, L)$, so system (27) might not be solvable. But by Claim 5, at most one cascade set, say $C_b(\delta)$, is reached across all actions. We can in this case construct a valid non-pivot mechanism as follows. Let the transfer for all active actions leading to $C_b(\delta)$ be 0. For all other active actions a , give the agent n a modified state contingent marginal contribution: $M'(a, \pi, \omega) = [\delta/(1 - \delta)](v(p(a, \pi, \xi), \omega) - u(b, \omega))$. \square

D.2 Mimicry with Two Actions: Proof of Corollary 5

For the sake of argument, consider action $a = 1$, with $p(1, \pi, \xi) \notin C(\delta)$ by assumption. First consider the case where also $p(2, \pi, \xi) \notin C(\delta)$. By equation (27),

$$t(1, 1) - t(1, 2) = \frac{M(1, \pi, L) - M(1, \pi, H)}{\psi(1, L) - \psi(1, H)}. \quad (28)$$

Now, $\psi(1, L) > \psi(1, H)$, as explained before (27). Thus, the fraction shares the sign of the numerator. In the definition of M in Proposition 4, $v(\pi, \rho)$ is a subtangent line of the value function at π . Then $M(1, \pi, L) - M(1, \pi, H) = \partial v(\pi, \rho)/\partial \rho - \partial v(p(1, \pi, \xi), \rho)/\partial \rho$. By the myopic action ordering, we have $p(1, \pi, \xi) < \pi$ and thus $\partial v(\pi, \rho)/\partial \rho - \partial v(p(1, \pi, \xi), \rho) \geq 0$, since the value function is convex. Then $t(1, 1) - t(1, 2) \geq 0$.

Finally, when $p(2, \pi, \xi) \in C(\delta)$ the logic is the same, substituting M in (28) by M' . \square

E THREE EXAMPLES

E.1 Calculations for the Binomial Signal Example in §4

We compute the value of the strategy in §4. For public beliefs $\pi \geq \pi_3$, the planner takes action 2 forever, and so $v(\pi) = 2\pi - 1$. Symmetrically, $v(\pi) = 1 - 2\pi$ for all $\pi \leq 1 - \pi_3$. By symmetry around $\pi = 1/2$, the value function is constant at $v(1/2)$ on $[1 - \pi_1, \pi_1]$.

Starting at any public belief $\pi \in [\pi_2, \pi_3)$, a high signal results in a continuation $\pi > \pi_3$

³⁴The intuitive pivot mechanism has an interesting implication we pursue below. Bergemann and Välimäki (2010) focus on flow marginal contribution. But in our case, each agent enters just once.




Figure 3: **Cascade Set Shrinks in δ .** This plots the value function $v(\pi)$ in the binomial signal example, for discount factor $\delta = 1/2$ and signals $\sigma_0 = 1/3 < 2/3 = \sigma_1$. The cascade set $[\pi_3, 1]$ is smaller than with discount factor $\delta = 0$, since $\pi_3 = 18/25 > 2/3 = \sigma_1$.

while a low signal leads to public beliefs in $[1 - \pi_1, \pi_1]$, where $v(\pi) \equiv 1/2$. By recursion,

$$v(\pi) = [\pi\sigma_1 + (1 - \pi)\sigma_0](2r(\pi, \sigma_1) - 1) + [\pi\sigma_0 + (1 - \pi)\sigma_1][(1 - \delta)(1 - 2r(\pi, \sigma_0)) + \delta v(1/2)].$$

Since $\pi_2 < \sigma_1 < \pi_3$, this formula for v obtains at $\pi = \sigma_1$. But by symmetry, $v(1/2) = (1 - \delta)(2\sigma_1 - 1) + \delta v(\sigma_1)$, namely, the myopic payoff plus the discounted continuation value. Solving these two equations, $v(1/2) = (2\sigma_1 - 1)/(1 - 2\delta^2\sigma_0\sigma_1)$.

The display formula for v at $\pi = \pi_3$ and $v(\pi_3) = 2\pi_3 - 1$ yield (9). Figure 3 plots v .

E.2 An Example of Actions in Non-Natural Order

To illustrate a non-natural action order asserted in §5, 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 , reflecting payoffs ± 1 when the action matches/mismatches the state. Choose a high discount factor $\delta = 0.95$.

Figure 4 depicts the numerically calculated private posterior belief threshold $\theta(\pi)$. For public beliefs $\pi \in (.3, .4) \subset (0, 3/7)$, 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), it gives no planner gain when

$$\begin{aligned} & \psi(2, \pi, \xi)(2p(2, \pi, \xi) - 1) + \psi(1, \pi, \xi)(1 - 2p(1, \pi, \xi)) \\ > & \psi(1, \pi, \xi)(2p(1, \pi, \xi) - 1) + \psi(2, \pi, \xi)(1 - 2p(2, \pi, \xi)) \end{aligned} \quad (29)$$




Figure 4: **Inverted Action Ordering.** The optimal posterior belief threshold $\theta(\pi)$.

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 (29) holds at low π , as the example shows, when the reversed order takes action 1 for a relatively large set of high signals. \square

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.

E.3 The Role of Posterior Monotonicity in Contrarianism

We show by an example that the posterior monotonicity property (Lemma 4) is necessary for contrarianism in Proposition 3 when the convex value function v can be chosen freely (see footnote 27). We use a version of the two-period professor-student example with $\delta = 1$ in §6.1 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 $g(\rho|\pi)$. This signal structure violates posterior monotonicity for some interval, say $[\hat{\theta}, 1]$. Thus,

$$p' \equiv \frac{\int_{\hat{\theta}}^1 \rho g(\rho|\pi') d\rho}{\int_{\hat{\theta}}^1 g(\rho|\pi') d\rho} > \frac{\int_{\hat{\theta}}^1 \rho g(\rho|\pi'') d\rho}{\int_{\hat{\theta}}^1 g(\rho|\pi'') d\rho} \equiv p''$$

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

Figure 5 illustrates the convex value function that we construct for the example. First




Figure 5: **Necessity Principle.** The student's value function for §E.3.

choose an arbitrary $\theta_{23} \in (p'', p')$. For any $\varepsilon > 0$, the convex function $\hat{v}(p|\varepsilon)$ consists of three linear segments $l_1, l_2, l_3(\varepsilon)$. Segments l_1, l_2 intersect at $\hat{\theta}$, while $l_2, l_3(\varepsilon)$ 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 $l_1, l_3(\varepsilon)$ 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 $\hat{\theta}, \theta_{12}, \theta_{23}(\varepsilon)$ 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 suboptimality 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 θ_{13} 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 intersect at $\hat{\theta}$, and therefore the first order condition 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 $l_3(\varepsilon)$. These cross at $\theta_{13}(\varepsilon)$, and so the first order condition fails at θ_{23} . If in $[\hat{\theta}, \theta_{23})$, the first order condition 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 $\hat{v}(\pi''|0)$. \square

References

- ALI, S. N., AND N. KARTIK (2012): “Herding with Collective Preferences,” *Economic Theory*, 51, 601–626.
- ANDERLINI, L., D. GERARDI, AND R. LAGUNOFF (2012): “Communication and Learning,” *Review of Economic Studies*, 79, 419–450.
- BANERJEE, A. V. (1992): “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 107, 797–817.
- BENVENISTE, L. M., AND J. A. SCHEINKMAN (1979): “On the Differentiability of the Value Function in Dynamic Models of Economics,” *Econometrica*, 47(3), 727–732.
- BERGEMANN, D., AND J. VÄLIMÄKI (2010): “An Essentially Complete Class of Decision Functions for Certain Standard Sequential Problems,” *Econometrica*, 78, 771–789.
- BERTSEKAS, D. (1987): *Dynamic Programming: Deterministic and Stochastic Models*. Prentice Hall, Englewood Cliffs, N.J.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): “A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades,” *Journal of Political Economy*, 100, 992–1026.
- BOSE, S., G. OROSEL, M. OTTAVIANI, AND L. VESTERLUND (2006): “Dynamic Monopoly Pricing and Herding,” *Rand Journal of Economics*, 37, 910–928.
- (2008): “Monopoly Pricing in the Binary Herding Model,” 37, 203–241.
- BREIMAN, L. (1968): *Probability*. Addison-Wesley, Reading, Mass.
- BRU, L., AND X. VIVES (2002): “Informational Externalities, Herding and Incentives,” *Journal of Institutional and Theoretical Economics*, 158, 91–105.
- CHAMLEY, C., AND D. GALE (1994): “Information Revelation and Strategic Delay in a Model of Investment,” *Econometrica*, 62, 1065–1085.
- CHE, Y.-K., AND J. HÖRNER (2016): “Optimal Design for Social Learning,” Columbia and Yale mimeo.
- COTTER, K. D., AND J.-H. PARK (2006): “Non-Concave Dynamic Programming,” *Economics Letters*, 90, 141–146.
- DEWATRIPONT, M., AND J. TIROLE (1994): *The Prudential Regulation of Banks*. MIT Press, Cambridge.
- 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 N. KIEFER (1988): “Controlling a Stochastic Process with Unknown Parameters,” *Econometrica*, 56, 1045–1064.
- ELY, J., A. FRANKEL, AND E. KAMENICA (2015): “Suspense and Surprise,” *Journal of Political Economy*, 123, 215–260.
- ELY, J. C. (2017): “Beeps,” *American Economic Review*, 107, 31–53.
- GILL, D., AND D. SGROI (2008): “Sequential Decisions with Tests,” *Games and Economic Behavior*, 63, 663–678.
- (2012): “The Optimal Choice of Pre-launch Reviewer,” *Journal of Economic Theory*, 147, 1247–1260.
- GITTINS, J. C. (1979): “Bandit Processes and Dynamical Allocation Indices,” *Journal of the Royal Statistical Society, Series B*, 41, 148–177.
- GLAZER, J., I. KREMER, AND M. PERRY (2015): “Crowd Learning without Herding: A Mechanism Design Approach,” Warwick mimeo.
- HERRERA, H., AND J. HÖRNER (2012): “A Necessary and Sufficient Condition for Information Cascades,” Columbia and Yale mimeo.
- HÖRNER, J., AND A. SKRZYPACZ (2016): “Selling Information,” *Journal of Political Economy*, 124, 1515–1562.
- JEGADEESH, N. (1990): “Evidence of Predictable Behavior of Security Returns,” *Journal of Finance*, 45, 881–898.
- KAMENICA, E., AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- KREMER, I., Y. MANSOUR, AND M. PERRY (2014): “Implementing the “Wisdom of the Crowd”,” *Journal of Political Economy*, 122, 988–1012.
- MARCH, C., AND A. ZIEGELMEYER (2016): “Altruistic Observational Learning,” CESifo working paper.
- MARSHALL, A. W., AND I. OLKIN (1979): *Inequalities: Theory of Majorization and Its Applications*. Academic Press, San Diego.
- MEDRANO, L. A., AND X. VIVES (2001): “Strategic Behavior and Price Discovery,” *Rand Journal of Economics*, 32, 221–248.

- MORRIS, S., AND M. YILDIZ (2016): “Crises: Equilibrium Shifts and Large Shocks,” SSRN working paper.
- PASTORINO, E., AND P. J. KEHOE (2011): “Individual and Market Experimentation with Multi-Armed Dependent Bandits,” FRB of Minneapolis mimeo.
- QUAH, J. K.-H., AND B. STRULOVICI (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.
- RAYO, L., AND I. SEGAL (2010): “Optimal Information Disclosure,” *Journal of Political Economy*, 118, 949–987.
- ROTHSCHILD, M. (1974): “A Two-Armed Bandit Theory of Market Pricing,” *Journal of Economic Theory*, 9, 185–202.
- SGROI, D. (2002): “Optimizing Information in the Herd: Guinea Pigs, Profits, and Welfare,” *Games and Economic Behavior*, 39, 137–166.
- SMITH, L., AND P. SØRENSEN (2000): “Pathological Outcomes of Observational Learning,” *Econometrica*, 68, 371–398.
- SMITH, L., AND J. TIAN (2016): “Informational Inertia,” Wisconsin mimeo.
- 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.