Observational Learning with Ordered States*

Navin Kartik[†] SangMok Lee[‡] Tianhao Liu[§] Daniel Rappoport[¶]

August 15, 2022

Abstract

When does society eventually learn the truth, or the correct action to take, via observational learning? In a general model of sequential learning over social networks, we develop the interplay of preferences satisfying single-crossing differences (SCD) and a new informational condition, directionally unbounded beliefs (DUB). For a wide class of network structures, SCD preferences and DUB information are a minimal pair of sufficient conditions for learning. Unlike "unbounded beliefs", which is the informational condition that characterizes learning for all preferences, DUB is broadly compatible with the monotone likelihood ratio property. Our analysis establishes a welfare lower bound under arbitrary preferences and information, which yields a general condition for learning. What drives learning is that a single agent must be able to rule out any wrong action, rather than (individually) take the correct action.

Keywords: social learning; herds; information cascades; single crossing; monotone likelihood ratio property; unbounded beliefs.

---

*This is a major revision of a 2021 working paper with the same title by Kartik, Lee, and Rappoport. That paper only studied a specific observational structure, the complete network. We thank Nageeb Ali, Marina Halac, Elliot Lipnowski, José Montiel Olea, Xiaosheng Mu, Evan Sadler, Lones Smith, and Peter Sørensen for helpful comments. We also received useful feedback from various seminar and conference audiences. César Barilla, John Cremin, and Zikai Xu provided excellent research assistance. Kartik gratefully acknowledges support from NSF Grant SES-2018948.

[†]Department of Economics, Columbia University. E-mail: nkartik@columbia.edu.

[‡]Department of Economics, Washington University in St. Louis. E-mail: sangmoklee@wustl.edu.

[§]Department of Economics, Columbia University. E-mail: tl3014@columbia.edu.

[¶]Booth School of Business, University of Chicago. E-mail: Daniel.Rappoport@chicagobooth.edu.

1. Introduction

This paper concerns the classic sequential observational or social learning model initiated by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992). There is an unknown payoff-relevant state (e.g., product quality). Each of many agents has homogeneous preferences over her own action and the state (e.g., all prefer products of higher quality). But agents act in sequence, each receiving her own private information about the state and observing some subset—not necessarily all—of her predecessors' actions. The central economic question is about asymptotic learning: do Bayesian agents eventually learn to take the correct action (e.g., will the highest quality product eventually prevail)?

Economists' answer to this fundamental question is that (asymptotic, Bayesian) social learning turns on whether private signals/beliefs are *unbounded*. More precisely, unbounded beliefs is the informational condition that characterizes learning for all preferences, so long as the observational network structure satisfies a mild condition referred to as "expanding observations" (Smith and Sørensen, 2000; Acemoglu, Dahleh, Lobel, and Ozdaglar, 2011; Arieli and Mueller-Frank, 2021). Roughly, unbounded beliefs says that given any (full-support) prior it should be possible for a single private signal, however unlikely it is, to make an agent arbitrarily close to certain about the true state. The canonical example is when there are two states, say $\omega \in \Omega = \{1, 2\}$, and there is *normal information*: agents' signals are drawn independently from normal distributions with mean ω and fixed variance. In this case very high signals make one arbitrarily convinced that $\omega = 2$, while very low signals make one arbitrarily convinced that $\omega = 1$.

There are two related limitations to this received wisdom. First, the result has only been established for two states; indeed, for Bayesian learning on general observational networks, we are not aware of any prior work that tackles multiple—i.e., more than two—states. This is a major shortcoming, as multiple states is the realistic case in any number of applications of observational learning (e.g., to product choice). Second, requiring unbounded beliefs with multiple states is extremely restrictive. To wit, if the state space $\Omega \subset \mathbb{R}$ has more than two states, then normal information no longer satisfies unbounded beliefs: given any full-support prior, there is an upper bound on how certain one can become about any non-extremal state based on observing one signal.[1] Is social learning doomed with multiple states for familiar information structures like normal information?

[1] So binary states is special because all states are extreme states; in general, there need not even be any extreme states (e.g., the states are all integers). There is nothing special about normal information violating unbounded beliefs. Under full support (i.e., no signal excludes any state), any information structure satisfying the widely-used monotone likelihood ratio property precludes unbounded beliefs with multiple states; see Subsection 4.2.

Our paper shows that the answer is no. It turns out that with multiple states, unbounded beliefs is not necessary for learning with standard preferences. In general, whether there is learning depends jointly on preferences and information. Our main substantive result is that for a canonical class of preferences considered in economic models—those with *single-crossing differences* (SCD)—learning obtains in general observational networks (satisfying the aforementioned expanding observations) when the information structure satisfies *directionally unbounded beliefs* (DUB). SCD is a familiar preference property (Milgrom and Shannon, 1994), which subsumes any supermodular utility function; we are not aware, however, of it having been studied in social learning. By contrast, DUB appears to be a new condition on information structures, although Milgrom (1979) is a precedent in the context of auction theory. Like SCD, DUB is formulated for an ordered state space. It requires that for any state ω and any prior that puts positive probability on ω, there exist both: (i) signals that make one arbitrarily certain that the state is at least ω; and (ii) signals that make one arbitrarily certain that the state is at most ω.[2] Crucially, no signal need make one arbitrarily certain about ω (unlike unbounded beliefs). For the normal information structure, requirements (i) and (ii) are met for any state by arbitrarily high and arbitrarily low signals, respectively. More generally, this is the case for any subexponential location-shift signal structure (Proposition 3 in Section 5); alternatively, given the monotone likelihood property, DUB holds under the basic requirement of "pairwise unbounded beliefs" (Remark 1 in Section 5).

Theorem 1 in Subsection 3.2 establishes that DUB information and SCD preferences are jointly sufficient for learning, or more precisely, adequate learning à la Aghion, Bolton, Harris, and Jullien (1991). For some intuition on their interplay, consider normal information again. There are preferences for which learning can fail because society may get stuck at some belief at which agents are taking an incorrect action, but only a strong signal about an intermediate state can change the action—alas no such signal is available. However, under SCD preferences, if a strong signal about some intermediate state ω would change the action, then so would either a signal that indicates the state is very likely to be at least ω or at most ω. Such signals are guaranteed in the normal example, and more generally by DUB information. We further establish in Section 4 that DUB information and SCD preferences are a *minimal* pair of sufficient conditions for adequate learning in the sense of Athey (2002): if one property fails, then there are primitives satisfying the other property such that *in*adequate learning obtains.

Our second contribution lies in developing an approach to tackle learning, and more

[2] This version of the definition assumes the state space is countable. Milgrom (1979) imposes only requirement (i), not (ii).

generally, (asymptotic) social welfare with multiple states in general observational networks. Theorem 2 in Subsection 3.3 is a methodological contribution that is both the backbone for proving Theorem 1, and also of independent interest. The theorem establishes a social welfare bound: roughly, for any preference and information (and given expanding observations), agents eventually obtain at least their *cascade utility*. Cascade utility is the minimum expected utility an agent can get from any Bayes-plausible distribution of stationary beliefs, i.e., beliefs at which it is optimal to ignore the private signal. Theorem 2 implies that adequate learning obtains when the cascade utility is the same as the utility obtained from taking the correct action in each state. This point underlies why DUB information is sufficient for adequate learning for all SCD preferences, and also yields that unbounded beliefs is sufficient for all preferences. More generally, Theorem 2 and its consequences, Lemma 2 and Proposition 5, should be valuable for future work interested in whether there is adequate learning for other combinations of preferences and information. Furthermore, the result also provides a quantification of welfare even when learning fails.

Related Literature. A number of papers on sequential (Bayesian) social learning only consider the complete observational network: each agent observes all her predecessors' actions. For that case, and focusing on binary states, Smith and Sørensen (2000) introduced the distinction between bounded and unbounded beliefs; they show that, given any non-trivial preference, there is learning if and only if beliefs are unbounded.[3] For the complete network but with multiple states, Arieli and Mueller-Frank (2021, Theorem 1) show that unbounded beliefs—which they call "totally unbounded beliefs"—is sufficient for learning, and also necessary if learning must obtain no matter society's preference.[4] Importantly, the approach of both Smith and Sørensen (2000) and Arieli and Mueller-Frank (2021) rests on the fact that in the complete observational network, the public belief—the belief held by an agent based on observing her predecessors' actions, before observing her private signal—is a martingale.

Gale and Kariv (2003) and Çelen and Kariv (2004) depart from the complete network, noting that martingale arguments now fail. Both these papers also depart from the canonical setting in other ways, however: in Gale and Kariv (2003) agents choose actions repeat-

[3] Our literature discussion should be understood as assuming a finite action space; it is well known since Lee (1993) that with an infinite action space, learning can obtain even with bounded beliefs. Variations on other dimensions can also generate learning with bounded beliefs, such as endogenous prices (Avery and Zemsky, 1998) or congestion costs (Eyster, Galeotti, Kartik, and Rabin, 2014), and heterogenous preference types (Goeree, Palfrey, and Rogers, 2006).

[4] The early work of Bikhchandani, Hirshleifer, and Welch (1992) allowed for multiple states, but they only identified failures of learning because they implicitly restricted attention to bounded beliefs; more precisely, they assumed finite signals with full-support distributions.

edly, while in Çelen and Kariv (2004) private signals are not conditionally independent given the true state. Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) provide a general treatment of observational networks in an otherwise classical setting. They only allow for binary states and binary actions, however. They introduce the condition of expanding observations on the observational network structure, showing that it is necessary for learning. Their Theorem 2 establishes that it is also sufficient for learning with unbounded beliefs. A key contribution of Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) is to use an *improvement principle* to deduce learning (under unbounded beliefs and expanding observations); this approach works even though martingale arguments fail. Lobel and Sadler (2015) introduce a notion of "information diffusion" and use the improvement principle to establish information diffusion even when learning fails; they also consider more general networks than Acemoglu, Dahleh, Lobel, and Ozdaglar (2011).

The analysis in both Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) and Lobel and Sadler (2015) relies on their binary-state binary-action structure.[5] Ours is the first paper to consider the canonical sequential social learning problem with general observational networks and general state and action spaces. At a methodological level, we develop a novel analysis based on continuity and compactness—rather than monotonicity or other properties that are specific to binary states or actions—that uncovers the fundamental logic underlying a general improvement principle. We use that to establish the cascade-utility welfare bound in any network satisfying expanding observations. As explained in Subsection 3.3, our cascade-utility bound coincides with Lobel and Sadler's (2015) information-diffusion bound in a binary-state binary-action model, but in general it is higher.

More importantly, our focus on multiple states and actions allows us to shed light on how preferences and information jointly determine whether there is social learning; in particular, we deduce that SCD preferences and DUB information are a minimal pair of sufficient conditions for learning. The interaction of preferences and information for learning has not received attention in the prior literature because of its focus on two states. An exception is Arieli and Mueller-Frank (2021, Theorem 3), but their result assumes a special utility function and is only for the complete network.

We leave to future research the speed of learning/welfare convergence, both under the SCD-DUB combination and beyond. For binary states and the complete network, Rosenberg and Vieille (2019) deduce the condition on the likelihood of extreme posteriors that determines whether learning is, in certain senses, efficient; they point out that their condition is violated by normal information. See Hann-Caruthers, Martynov, and Tamuz (2018)

[5] Banerjee and Fudenberg (2004) and Smith and Sørensen (2020) consider "unordered" random sampling models that also only allow for binary states and actions.

as well.

Lastly, we note that there is a large literature on non-Bayesian social learning, surveyed by Golub and Sadler (2016). There has also been recent interest in (mis)learning among misspecified Bayesian agents; see, for example, Frick, Iijima, and Ishii (2020, 2021) and Bohren and Hauser (2021).

2. Model

In the main text of the paper, we focus on a model that simplifies some technical issues. See Appendix A for a more general setting.

There is a countable (finite or infinite) set of states $\Omega \subset \mathbb{R}$, and standard Borel spaces of actions A and signals S. An information or signal structure is given by a collection of probability measures over S, one for each state, denoted by $F(\cdot|\omega)$. Assume that for any ω and ω', $F(\cdot|\omega)$ and $F(\cdot|\omega')$ are mutually absolutely continuous. It follows that each $F(\cdot|\omega)$ has a density $f(\cdot|\omega)$; more precisely, this is the Radon-Nikodym derivative of $F(\cdot|\omega)$ with respect some reference measure that is mutually absolutely continuous with every $F(\cdot|\omega')$. Without further loss of generality we assume $f(s|\omega) > 0$, so that no signal excludes any state.

The game. At the outset, date 0, a state ω is drawn from a common prior probability mass function $\mu_0 \in \Delta\Omega$.[6] An infinite sequence of agents, indexed by $n = 1, 2, \ldots$, then sequentially select actions. An agent n observes both a private signal s_n drawn independently from $f(\cdot|\omega)$ and the actions of some subset of her predecessors $B_n \subseteq \{1, 2, \ldots, n-1\}$, and then chooses her action $a_n \in A$. No agent observes either the state or any of her predecessors' signals. Each observational neighborhood B_n is stochastically generated according to a probability distribution Q_n over all subsets of $\{1, 2, \ldots, n-1\}$, assumed to be independent across n, independent of the state ω, and independent of any private signals. The distributions $(Q_n)_{n\in\mathbb{N}}$ constitute the observational network structure and are common knowledge, but the realization of each neighborhood B_n is the private information of agent n.

Agent n's information set thus consists of her signal s_n, neighborhood B_n, and the actions chosen by the neighbors $(a_k)_{k\in B_n}$. Let $\mathcal{I}_n$ denote the set of all possible information sets for agent n. A strategy for agent n is a (measurable) function $\sigma_n : \mathcal{I}_n \to \Delta A$.

All agents are expected utility maximizers and share the same preferences over their actions: every agent n's state-dependent preferences are represented by the von-Neumann-

[6] For any measurable space X, ΔX denotes the set of probability measures over X.

Morgenstern utility function $u : A \times \Omega \to \mathbb{R}$. To ensure expected utility is well defined for all probability distributions, we assume utility is bounded: there is $\overline{u} \geq 0$ such that $|u(\cdot,\cdot)| \leq \overline{u}$.

We study the Bayes Nash equilibria—hereafter simply equilibria—of this game. An equilibrium exists if for every belief there is an optimal action, which we will implicitly assume to be the case.[7]

Adequate learning. The full-information expected utility given a belief μ is the expected utility under that belief if the state will be revealed before an action is chosen:

$$u^*(\mu) := \sum_{\omega \in \Omega} \max_{a \in A} u(a,\omega)\mu(\omega).$$

Given a prior μ_0 and a strategy profile σ, agent n's utility u_n is a random variable. Let $\mathbb{E}_{\sigma,\mu_0}[u_n]$ be agent n's ex-ante expected utility. We say there is *adequate learning* if for every prior μ_0 and every equilibrium strategy profile σ, $\mathbb{E}_{\sigma,\mu_0}[u_n] \to u^*(\mu_0)$. In words, adequate learning requires that given any prior and equilibrium, no matter which state is realized, eventually agents take actions that are arbitrarily close to optimal in that state. We say there is *inadequate learning* if adequate learning fails.[8]

Expanding observations. As observed by Acemoglu, Dahleh, Lobel, and Ozdaglar (2011), a necessary condition for adequate learning is that the network structure have *expanding observations*:

$$\forall K \in \mathbb{N} : \lim_{n\to\infty} Q_n\left(B_n \subseteq \{1,\ldots,K\}\right) = 0.$$

A failure of expanding observations means that for some $K \in \mathbb{N}$, there is an infinite number of agents each of whom, with probability uniformly bounded away from 0, observes the actions of only at most the first K agents. Plainly, that would preclude adequate learning because it would imply that an infinite number of agents, with probability uniformly bounded away from 0, cannot do better than choosing their action based on only $K+1$ signals.

[7] Existence of optimal actions is assured under standard assumptions, e.g., A is finite, or more generally, A is compact and $u(\cdot,\cdot)$ is suitably continuous. We also note that as there are no direct payoff externalities, strategic interaction is minimal: any σ_n affects other agents only insofar as affecting how n's successors update about signal s_n from the observation of action a_n. We could just as well adopt (weak) Perfect Bayesian equilibrium or refinements.

[8] That we deem learning to be inadequate if there is some equilibrium in which learning fails, rather than in every equilibrium, is not essential, given that there is no strategic interaction (cf. fn. 7). On the other hand, whether learning fails at every prior rather than only at some priors is important, and addressed in Subsection 5.2 by Proposition 4.

Accordingly, we assume expanding observations. Here are a few leading examples of network structures with expanding observations: (i) the classic complete network in which each agent's neighborhood is all her predecessors (formally, $Q_n(B_n = \{1, \dots, n-1\}) = 1$); (ii) each agent only observes her immediate predecessor ($Q_n(B_n = \{n-1\}) = 1$); and (iii) each agent observes a random predecessor ($Q_n(B_n = \{k\}) = 1/(n-1)$ for all $k \in \{1, \dots, n-1\}$).[9]

3. Main Results

3.1. SCD & DUB

We now introduce the two conditions that are central to our main result, Theorem 1 below.

3.1.1. Single-Crossing Differences

A function $h : \mathbb{R} \to \mathbb{R}$ is *single crossing* if either: (i) for all $x < x'$, $h(x) > 0 \implies h(x') \geq 0$; or (ii) for all $x < x'$, $h(x) < 0 \implies h(x') \leq 0$. That is, a single-crossing function switches sign between strictly positive and strictly negative at most once. We are interested in the following class of preferences.

Definition 1. Preferences represented by $u : A \times \Omega \to \mathbb{R}$ have *single-crossing differences* (SCD) if for all a and a', the difference $u(a, \cdot) - u(a', \cdot)$ is single crossing.

SCD is an ordinal property closely related to a notion in Milgrom and Shannon (1994), but, following Kartik, Lee, and Rappoport (2019), the formulation is without an order on A. Ignoring indifferences, SCD requires that the preference over any pair of actions can only flip once as the state changes monotonically. If A is ordered, then SCD is implied by Milgrom and Shannon's (1994) notion or even a weaker one in Athey (2001).[10] SCD is thus a property widely satisfied by economic models; in particular, it is assured by supermodularity of u.

[9] In our model, for every action that an agent observes, she knows the identity of the predecessor who took that action. So in the random-predecessor case, even though it is random which predecessor an agent observes, the predecessor's identity is known. Our methods extend to cover some forms of random sampling in which an agent does not know which predecessors' actions she observes.

[10] Our treatment of indifferences is more permissive than Milgrom and Shannon (1994) and Kartik, Lee, and Rappoport (2019). A function $h : \{1, 2, 3\} \to \mathbb{R}$ given by $h(1) = 0$, $h(2) = 1$, and $h(3) = 0$ is single crossing per our definition but not according to the notions employed by Milgrom and Shannon (1994) and Kartik, Lee, and Rappoport (2019). SCD is, in fact, equivalent to saying that there exists some order on A with respect to which Athey's (2001) notion of "weak single-crossing property of incremental returns" is satisfied.

3.1.2. Directionally Unbounded Beliefs

Given any state ω and belief μ, let $\mu(\omega|s)$ be the posterior after signal s, and $\mu(\omega|\cdot)$ the corresponding random variable. We say that ω is *distinguishable* from $\Omega' \subseteq \Omega\backslash\{\omega\}$ if for every $\mu \in \Delta(\Omega' \cup \omega)$ with $\mu(\omega) > 0$, $\operatorname{Supp}\mu(\omega|\cdot) \ni 1$. That is, starting from any belief that puts positive probability on ω, there are signals that make ω arbitrarily more likely relative to Ω'.

Distinguishability of each state from all other states is the condition of *unbounded beliefs*; this is termed "totally unbounded beliefs" by Arieli and Mueller-Frank (2021) and is the multi-state extension of the two-state notion introduced by Smith and Sørensen (2000). But unbounded beliefs is highly restrictive with more than two states; specifically, any signal structure satisfying a ubiquitous property in information economics—the monotone likelihood ratio property (MLRP)—violates unbounded beliefs (see Subsection 4.2 below).[11] Accordingly, we introduce the following weaker condition.

Definition 2. There is *directionally unbounded beliefs* (DUB) if every ω is distinguishable from $\{\omega' : \omega' < \omega\}$ and also from $\{\omega' : \omega' > \omega\}$.

Crucially, DUB does not require or assure that any state ω is distinguishable from all other states—not even any subset of states containing both some higher and some lower state than ω. What DUB does assure, and will be crucial in our analysis, is that for any ω and prior μ with $\mu(\omega) > 0$, there are signals that make the posterior on the lower set $\{\omega' : \omega' < \omega\}$ arbitrarily small, and analogously for the upper set $\{\omega' : \omega' > \omega\}$.

Lemma 1. *DUB holds if for every ω and $\varepsilon > 0$, there exist positive probability sets of signals $\overline{S}, \underline{S}$ such that:*

1. $\forall \omega' < \omega, \forall s \in \overline{S}$: $\frac{f(s|\omega')}{f(s|\omega)} < \varepsilon$;
2. $\forall \omega' > \omega, \forall s \in \underline{S}$: $\frac{f(s|\omega')}{f(s|\omega)} < \varepsilon$.

Conversely, this condition is implied by DUB when Ω is finite.

Lemma 1 provides an intuitive sufficient condition for DUB that is also necessary when Ω is finite.[12] To interpret the condition, fix any state ω. Condition 1 of the lemma requires

[11] A signal density $f(s|\omega)$ satisfies the MLRP if the signal space is $S \subset \mathbb{R}$ and $\forall s' > s$ and $\forall \omega' > \omega$, $\frac{f(s|\omega')}{f(s|\omega)} \leq \frac{f(s'|\omega')}{f(s'|\omega)}$.

[12] Milgrom (1979, Theorem 2) makes essentially the same observation, albeit when each state has only a finite number of lower states; in the context of auctions, he is only concerned with distinguishing each state from its lower set (not also its upper set). As he mentions, it is possible that any/all states are distinguishable from their lower sets but not their upper sets, or vice-versa, or they can be distinguishable from both or neither.

signals that are arbitrarily more likely in ω relative to every $\omega' < \omega$. We stress that while the set $\overline{S}$ can depend on ω, it cannot depend on the lower state $\omega' < \omega$. This ensures that starting from any prior with support $\{\omega' : \omega' \leq \omega\}$, there are signals that lead to posteriors arbitrarily near certainty on ω; that is, ω is distinguishable from all lower states. Condition 2 of the lemma analogously implies distinguishability of ω from all higher states.

A leading example of DUB information is when signals are normally distributed on $\mathbb{R}$ with mean ω and fixed variance; we hereafter refer to this specification as the *normal signal structure* or just *normal information*. For $\Omega \subset \mathbb{Z}$, one can see that DUB is satisfied here by noting that condition 1 of Lemma 1 is met because $\frac{f(s|\omega-1)}{f(s|\omega)} \to 0$ as $s \to \infty$ and for any $\omega' < \omega$, $\frac{f(s|\omega')}{f(s|\omega)} \leq \frac{f(s|\omega-1)}{f(s|\omega)}$; analogously, taking $s \to -\infty$ verifies condition 2 of the lemma.[13] Notably, however, no state can be simultaneously distinguished from both some lower one and some higher one. Figure 1a provides a graphical depiction. The solid blue curve therein is the set of posteriors obtained under normal information when $\Omega = \{1, 2, 3\}$ and the prior is μ. Posteriors are bounded away from certainty on state 2. Yet they can become arbitrarily certain about the relative probability of state 2 to 1 (as the curve approaches state 3) and, separately, about the relative probability of state 2 to 3 (as the curve approaches state 1): this is seen by projecting both the prior and the set of posteriors onto the relevant edges; for example, when the prior is μ' (the projection of μ onto the 2–3 edge) the set of posteriors is the entire 2–3 edge (the projection of the blue curve onto the 2–3 edge).

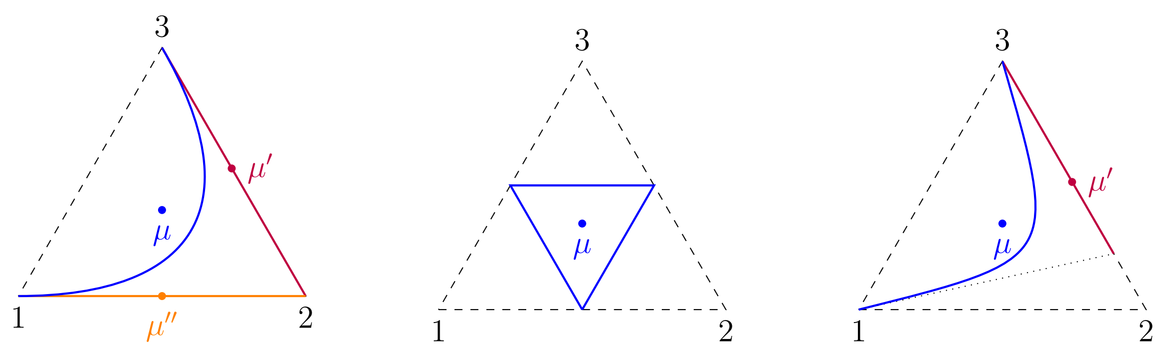

(a) DUB (b) Not DUB (c) Not DUB

Figure 1: Illustration of DUB. For each belief, posteriors are shown in the corresponding color.

Figure 1b and Figure 1c depict failures of DUB.[14] The information structure corresponding to the triangular set of posteriors in Figure 1b has the property that no signal can distin-

[13] This argument applies with only notational changes if $\Omega \not\subset \mathbb{Z}$, so long as every state has a maximum strictly lower state (if there are any strictly lower states) and a minimum strictly higher state (if there are any strictly higher states). Even absent this property (e.g., $\Omega = \mathbb{Q}$), normal information satisfies DUB by Proposition 3 in Subsection 5.1.

[14] As is well known (e.g., Aumann and Maschler, 1995; Kamenica and Gentzkow, 2011), a set of posteriors can be induced by some information structure if the prior is in the (relative) interior of the set's convex hull.

guish state 3 from states 2 and 1 simultaneously, even though a signal leading to a posterior close to the 1–3 edge (resp., the 2–3 edge) distinguishes state 3 from 2 (resp., from 1). By contrast, in Figure 1c states 1 and 3 are distinguishable from their complements, as the posteriors under the full-support prior μ, given by the solid blue curve, approach the relevant vertices. However, state 2 is not distinguishable from state 3: the posteriors for the prior μ' (the projection of μ onto the 2–3 edge), given by the solid red line (the projection of the blue curve onto the 2–3 edge), do not approach state 2's vertex. The reason is that unlike the case of normal information in Figure 1a, the blue curve in Figure 1c is not tangent to the 1–2 edge at the vertex 1.

3.2. Sufficient Conditions for Learning

Our main result is:

Theorem 1. *If preferences have SCD and the information structure has DUB, then there is adequate learning.*

Theorem 1 decouples conditions on preferences and information that are sufficient—indeed, minimally sufficient, as we will see in Section 4—for adequate learning. The theorem follows from a general characterization of learning that is jointly over preferences and information. To state that characterization, we require two concepts concerning the value of information.

Let $c : \Delta\Omega \rightrightarrows A$ be given by $c(\mu) := \arg\max_{a\in A} \mathbb{E}_\mu[u(a, \omega)]$. So $c(\mu)$ is the set of optimal actions under belief μ. Abusing notation, for a degenerate belief on state ω we write $c(\omega)$. Denoting the posterior after signal s when starting from belief μ by $\mu(\cdot|s)$, we say that belief μ is *stationary* if there is $a \in c(\mu)$ such that $a \in c(\mu(\cdot|s))$ for μ-a.e. signals s. We say that belief μ has *adequate knowledge* if there is $a \in c(\mu)$ such that $a \in c(\omega)$ for all $\omega \in \operatorname{Supp} \mu$. In words, a belief is stationary if an agent holding that belief does not benefit from observing a signal from the given information structure; for, there is some action that is optimal across all signal realizations (except perhaps for a set of measure zero) from the given information structure. On the other hand, a belief has adequate knowledge if the agent would not benefit from observing a signal from any information structure; for, there is some action that is optimal across all states she ascribes positive probability to.

Lemma 2. *There is adequate learning if and only if all stationary beliefs have adequate knowledge.*

The "only if" direction of this result is straightforward because our notion of learning considers all priors: if an inadequate-knowledge belief is stationary, then when that belief is the prior, society is stuck with all agents taking the prior-optimal action even though that is not optimal in some states. More important and subtle is the lemma's "if" direction. Although it has precedents in the literature for specific environments, the logic in the current setting of general observational networks is fundamentally new. We elaborate on Lemma 2 in the next subsection, focusing instead now on how it leads to Theorem 1.

Lemma 2 reduces the question of learning to whether the set of stationary beliefs coincides with the set of adequate-knowledge beliefs. In general, stationary beliefs can have inadequate knowledge because of a mismatch between preferences and information: society's information structure may not provide enough information about the states that agents would benefit from learning about.

However, SCD and DUB together preclude such a mismatch. Here is why, with Figure 2 illustrating. Suppose $\Omega = \mathbb{Z}$ and a belief μ has inadequate knowledge, so that $c(\mu) \neq c(\omega^*)$ for some state $\omega^* \in \operatorname{Supp} \mu$. (For simplicity, suppose $c(\mu)$ and $c(\omega^*)$ are singletons.) In the figure, $\omega^* = 0$ and belief μ is shown by the blue bars. SCD implies that not only is $c(\omega^*)$ strictly preferred to $c(\mu)$ at ω^*, but $c(\omega^*)$ is also weakly preferred either at all higher states or at all lower states. Suppose the former case, as in the figure. DUB ensures that the posterior likelihood ratio of ω^* to its lower set $\{\omega' : \omega' < \omega^*\}$ can be arbitrarily high with positive probability. This implies that with positive probability, the posterior on $\{\omega' : \omega' \geq \omega^*\}$ is arbitrarily close to one; the figure depicts such a posterior using the red bars. Hence, with positive probability, an agent's signal will lead her to strictly prefer $c(\omega^*)$ to $c(\mu)$, which means μ is not stationary. It follows that under SCD and DUB, all stationary beliefs have adequate knowledge.

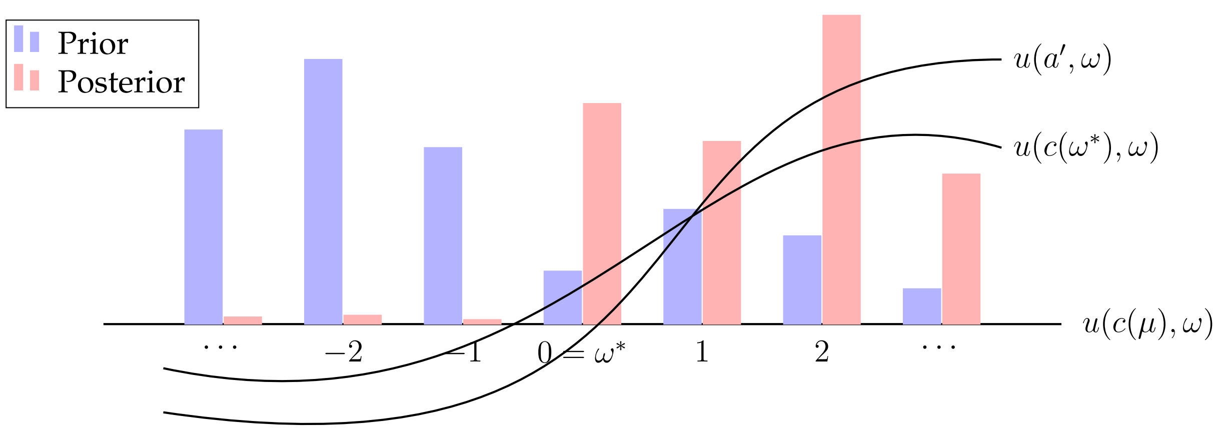

Figure 2: Illustration of Theorem 1. Blue bars represent the belief μ; red bars represent the posterior after observing a signal that provides near-certainty about ω^* relative to $\{\omega : \omega < \omega^*\}$. Black curves represent the utilities from action a' and $c(\omega^*)$. The utility of action $c(\mu)$ is normalized to 0.

Let us highlight that our argument in the previous paragraph—and more generally, the logic underlying Lemma 2—only requires that every wrong action (i.e., one that is suboptimal at the true state) can always be ruled out (i.e., starting from any belief with inadequate knowledge, with positive probability some signal will make that action suboptimal). But there is no guarantee that there is a signal leading to the correct action (i.e., one that is optimal at the true state). For example, in Figure 2, it could be that $c(\omega^*)$ is the correct action, but under the red-barred posterior showing that belief μ is not stationary, the optimal action is a' rather than $c(\omega^*)$; indeed, given belief μ there is no guarantee of any signal inducing an agent to take action $c(\omega^*)$.

By contrast, social learning has traditionally focused on the stronger property that it is always possible to induce an agent to take the correct action; this has lead to the emphasis on unbounded beliefs. When there are two states and finite actions, always being able to rule out a wrong action and always being able to induce the correct action are equivalent, as they both simply reduce to unbounded beliefs. But not so more generally. To see this sharply, consider $\Omega = A = \mathbb{Z}$, normal information, and $c(\omega) = \omega$ for all ω. For any given state and any full-support prior, the posterior probability is bounded away from 1; indeed, this bound is arbitrarily small for suitable priors. This means that there is no action that will be taken with positive probability at every full-support prior.[15] So, unlike in the two-state case, one cannot argue that the correct action always has a positive probability of overturning an incorrect herd. Rather, the logic behind Theorem 1 is that given SCD preferences, DUB information ensures that any incorrect herd will be overturned because some agent will learn that the correct action is superior (because of near-certainty on the corresponding lower/upper states), even though the requisite signal may lead the agent to take yet some other action.

3.3. Lemma 2 and a General Welfare Bound

We now return to Lemma 2. The lemma is best understood as a corollary of a result that provides a welfare bound regardless of whether there is learning. Stating that result requires some notation. Abusing notation, let

$$u(\mu) := \max_{a \in A} \sum_{\omega} u(a, \omega)\mu(\omega)$$

[15] This point could have also been made with a three state example, with the caveat that it only applies to non-extreme actions. In general, given normal information and distinct optimal actions in all states, no non-extreme action will be taken with positive probability at every full-support prior.

be the maximal utility an agent can get under belief μ, and let

$$I(\mu) := \left(\sum_{\omega \in \Omega} \int_S u(\mu_s) \, \mathrm{d}F(s|\omega)\mu(\omega) \right) - u(\mu)$$

be the utility improvement from observing a private signal at belief μ. Observe that $I(\mu) = 0$ for any stationary belief μ. We write $\Phi^{BP} \subset \Delta\Delta\Omega$ to denote the set of Bayes-plausible distributions of beliefs (i.e., $\mathbb{E}_\varphi[\mu] = \mu_0 \iff \varphi \in \Phi^{BP}$). Again abusing notation, we write $u(\varphi) := \mathbb{E}_\varphi[u(\mu)]$ for the utility of an agent under the distribution of beliefs φ, and analogously write $I(\varphi) := \mathbb{E}_\varphi[I(\mu)]$. It follows that

$$\Phi^S := \left\{ \varphi \in \Phi^{BP} : I(\varphi) = 0 \right\}$$

is the set of Bayes-plausible belief distributions that are supported on the set of stationary beliefs. (We have suppressed the dependence of Φ^{BP} and Φ^S on the prior μ_0.)

Building on a notion mentioned by Lobel and Sadler (2015), we can now define the *cascade utility* level as

$$u_*(\mu_0) := \inf_{\varphi \in \Phi^S} u(\varphi).$$

In words, $u_*(\mu_0)$ is the lowest utility level that an agent can get if her distribution of beliefs is supported on stationary beliefs. Our welfare bound, stated loosely, is that eventually all agents are assured a utility level of at least $u_*(\mu_0)$. More precisely:

Theorem 2. *In any equilibrium σ, $\liminf \mathbb{E}_{\sigma,\mu_0}[u_n] \geq u_*(\mu_0)$.*

We highlight that Theorem 2 does not presume SCD or DUB, nor does it require any order structure on states. The only substantive requirement is our maintained assumption that the observational network structure satisfies expanding observations. It is straightforward to see how Theorem 2 implies the "if" direction of Lemma 2. When all stationary beliefs have adequate knowledge, it holds for any distribution $\varphi \in \Phi^S$ that almost surely a correct action is taken. Thus, $u_*(\mu_0) = u^*(\mu_0)$, and we have adequate learning.

The argument behind Theorem 2 relies fundamentally on certain compactness and continuity. First, we establish in Lemma 4 of the appendix that Φ^{BP} is compact when the space of belief distributions is endowed with the Prohorov metric (which metrizes the weak topology). This is a consequence of Bayes-plausibility. Intuitively, for the case of countable states that our main text focuses on, although the prior μ_0 can be supported on an infinite set, it must concentrate an arbitrarily large mass on only finitely many states. More generally, for any $\delta > 0$, there exists a compact subset $\Omega' \subseteq \Omega$ such that $\mu_0(\Omega') \geq 1 - \delta$.[16] So any $\varphi \in \Phi^{BP}$ must put arbitrarily large probability on beliefs that put arbitrarily large probability on the compact set Ω' by Bayes plausibility, which yields compactness of Φ^{BP}. Next, owing to expected utility being continuous in beliefs, the improvement function $I(\varphi)$ is also continuous (Lemma 5 in the appendix), and thus uniformly continuous on Φ^{BP}. Consequently, $I(\varphi)$ achieves a minimum on any closed, hence compact, subset of Φ^{BP}.

Now consider any ε-neighborhood of the set of Bayes-plausible distributions supported on stationary beliefs, call it $(\Phi^S)^\varepsilon$. If an agent's belief distribution is in $(\Phi^S)^\varepsilon$, then her ex-ante expected utility is at least close to u_* since $I(\varphi) = 0$ on Φ^S and $I(\varphi)$ is uniformly continuous. On the other hand, if the belief distribution is not in $(\Phi^S)^\varepsilon$, then there is some strictly positive minimum utility improvement that the agent obtains (as the complement of $(\Phi^S)^\varepsilon$ is closed).

We can then apply an argument akin to the improvement principle of Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) and Lobel and Sadler (2015). Specifically, expanding observations guarantees that we can partition society into "generations" such that an agent in one generation observes, with large probability, a predecessor who is in either the previous generation or the current generation. Consider an agent's ex-ante expected utility. Either it is at least close to u_* (when the agent's belief distribution is in $(\Phi^S)^\varepsilon$), or the agent can improve her utility by a fixed amount compared to the agents she observes (because without a private signal she can simply mimic the highest-utility predecessor she observes). Thus, the lowest ex-ante expected utility in each generation increases by a fixed amount until it becomes at least close to u_*. Since ε was arbitrary, it follows that eventually all agents' utility must be higher than a level arbitrarily close to u_*, which is the conclusion of Theorem 2.

Although previous authors have deduced versions of Lemma 2 and Theorem 2 in special environments, what allows us to establish these two general results is our novel proof methodology. We highlight two distinctions with Lobel and Sadler (2015, Theorem 1), which is the most related existing result.[17] They consider a binary-state binary-action model. In that setting, they establish a welfare bound of "diffusion utility", which is the utility obtained by a hypothetical agent who observes an information structure that contains only the strongest signals (i.e., an "expert agent", in their terminology). Our cascade

[16] This statement applies for our general setting discussed in Appendix A, where we only require Ω to be sigma-compact, i.e., it is a countable union of compact sets. That covers, for example, the case of $\Omega = \mathbb{R}$.

[17] In the complete network, where each agent observes every predecessor's action, Theorem 2 can be proved by a martingale argument, as done by Arieli and Mueller-Frank (2021, Lemma 1). They argue that the asymptotic "public belief" has to be stationary. However, this approach is inoperable in general networks because the public belief need not be a martingale and its convergence is not guaranteed.

utility is more fundamentally tied to when learning stops, as it is defined using stationary beliefs. It is not hard to see that in general, no matter the number of states or actions, cascade utility is always at least as high as (the natural extension of) diffusion utility.[18] As Lobel and Sadler (2015) note, cascade utility and diffusion utility coincide in their binary-state binary-action model; but we note that in general, the former can be strictly higher.[19] Methodologically, Lobel and Sadler's (2015) argument for a minimum improvement follows Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) and relies on certain monotonicity that owes to their binary-binary setting. Our argument, by contrast, uses continuity of the improvement function and compactness of the set of Bayes-plausible beliefs.

4. Minimal-Sufficiency of SCD and DUB

In this section, we establish a sense in which SCD and DUB are a minimal pair of sufficient conditions for adequate learning: if one fails, then there are primitives satisfying the other such that there is inadequate learning.

4.1. Necessity of DUB

Proposition 1. *If the information structure does not have DUB, then there is inadequate learning for some SCD preferences.*

This result identifies DUB as the right informational condition to guarantee learning for all SCD preferences. The argument is as follows. Assume DUB fails: say, some state ω^* cannot be distinguished from all lower states (it is analogous if "lower" is replaced by "upper"). Consider a single-crossing utility difference between two actions a^* and $\underline{a}$ given by $\varepsilon > 0$ in states $\omega \geq \omega^*$, and -1 in states $\omega < \omega^*$. Any other actions can be assumed to be dominated and ignored. The failure of DUB implies that for any prior supported on $\omega \leq \omega^*$, the posterior on ω^* is bounded away from 1. But for any such prior, a^* is chosen only if the posterior on ω^* is greater than $1/(1+\varepsilon)$. It follows that for small enough ε, any

[18] Lobel and Sadler's (2015) definition of diffusion utility is tailored to their binary-binary model. In general, we can define it as the highest utility an agent can obtain from any Bayes-plausible belief distribution that is supported on the set of feasible posteriors (i.e., those available in the given information structure); call the corresponding signal structure the expert signal structure. Since expert signals are a subset of the original signals, the definition of stationary beliefs implies that observing a signal from the expert signal structure cannot improve an agent's utility when starting from a stationary belief. Hence, cascade utility is no lower than diffusion utility.

[19] This is true even with just two states and three actions. Alternatively, consider $\Omega = \{0, 1\}$, $A = [0, 1]$, $u(a, \omega) = -(a - \omega)^2$, and the complete network. It is well known that any nontrivial information structure leads to learning here, with the stationary beliefs being just 0 and 1. So the cascade utility is the full-information utility of 0, whereas diffusion utility will be strictly lower absent unbounded beliefs.

such prior will be stationary even though it has inadequate knowledge. By Lemma 2, there is inadequate learning.

4.2. Necessity of SCD

Given DUB, SCD is not necessary for adequate learning; after all, a violation of SCD can occur over a pair of actions that are both dominated.[20] However, we focus on SCD preferences because SCD *is* necessary given DUB when we vary choice sets, as is familiar in the comparative-statics literature. To make the point precise, say that there is *(in)adequate learning for a choice set* $\tilde{A} \subset A$ in the natural way, i.e., by replacing A with $\tilde{A}$ in all relevant definitions.

Proposition 2. *If a preference does not have SCD, then under any information structure with MLRP (including those with DUB), there is inadequate learning for some (binary) choice set.*

Recall that our leading example of DUB information is normal information, which satisfies MLRP. The intuition for Proposition 2 is simple. Consider a utility function that violates SCD. That is, there are actions $a_1, a_2 \in A$ and states $\omega_1 < \omega_2 < \omega_3$ such that $u(a_1, \omega_i) - u(a_2, \omega_i)$ is strictly positive for $i \in \{1, 3\}$ and strictly negative for $i = 2$. MLRP implies that signals' informativeness about state ω_2 relative to $\{\omega_1, \omega_3\}$ is bounded.[21] Hence, given choice set $\tilde{A} = \{a_1, a_2\}$, any belief that is supported on $\{\omega_1, \omega_2, \omega_3\}$ and puts sufficiently small weight on ω_2 relative to each of ω_1 and ω_3 is stationary. Since such priors have inadequate knowledge, Lemma 2 implies that there is inadequate learning.

Kartik, Lee, and Rappoport (2019, Theorem 1.1) establish that SCD is equivalent to "interval choice" across all choice sets: for every $\tilde{A} \subset A$, if action $a \in \tilde{A}$ is uniquely optimal in state ω_1 and in state $\omega_2 > \omega_1$, then $a \in \tilde{A}$ is also optimal in every state $\omega \in (\omega_1, \omega_2)$. But for a fixed choice set, even if there is a distinct optimal action in every state, learning can fail under DUB when SCD is violated. The following example illustrates.

Example 1. Let $\Omega = \mathbb{Z}$ and $A = \mathbb{Z} \cup \{a^*\}$. In any state ω, the utility from any integer action a is given by quadratic loss, $u(a, \omega) = -(a - \omega)^2$, whereas the action a^* is a "safe action", $u(a^*, \omega) = -\varepsilon$ for a small constant $\varepsilon > 0$.[22] So any action ω is uniquely optimal in state ω but worse than the safe action a^* in every other state. Plainly, SCD is violated.

[20] In fact, one can show that given DUB, Quah and Strulovici's (2009) interval dominance order with respect to some order on the action space ensures adequate learning under some regularity conditions (e.g., a finite action set).

[21] This uses our maintained assumption that the signal density is positive everywhere, which rules out examples like fully-revealing information.

[22] Strictly speaking, given $\Omega = \mathbb{Z}$, the quadratic loss utility function violates our maintained assumption of bounded utility, but we ignore this to keep the example succinct.

Consider normal information. There are full-support priors such that the posterior probability of any state is uniformly bounded away from one across all the signals and states.[23] For any such prior, for small enough $\varepsilon > 0$, the safe action a^* is optimal after every signal. In other words, any such prior is stationary but has inadequate knowledge, and so Lemma 2 implies that there is inadequate learning. $\square$

5. Discussion

5.1. DUB in Familiar Information Structures

Location families. For tractability, scholars often assume signal structures described by location-shift families of distributions—normal information being one example. It is of interest to know when DUB holds in these semi-parametric structures.

Formally, we say that the signal structure is a *location-shift signal structure* if $S = \mathbb{R}$ and there is a density $g : \mathbb{R} \to \mathbb{R}_+$, referred to as the *standard density*, such that $f(s|\omega) = g(s-\omega)$. We say that g is *strictly subexponential* if there is some $p > 1$ ($p \in \mathbb{R}$) such that for all s large enough in absolute value, $g(s) \leq \exp(-|s|^p)$. Normal information satisfies this condition with any $p \in (1,2)$.[24] Intuitively, a strictly subexponential density has a thin tail—it eventually decays faster than the exponential density.

Proposition 3. *For a location-shift signal structure with a uniformly continuous standard density g, DUB holds if g is strictly subexponential.*

The exponent p being strictly larger than one in the definition of strictly subexponential is essential for the result. To see that, consider the Laplace or double exponential standard density $g(s) = (1/2)\exp(-|s|)$. This density is not strictly subexponential, and indeed DUB fails: for any $\omega' < \omega < s$, $f(s|\omega')/f(s|\omega) = g(s-\omega')/g(s-\omega) = \exp(\omega' - \omega)$ is in-

[23] Take any prior μ such that for some $c > 0$, $\min\left\{\frac{\mu(n-1)}{\mu(n)}, \frac{\mu(n+1)}{\mu(n)}\right\} > c$ for all n (e.g., a double-sided geometric distribution). Denoting the posterior after signal s by μ_s, the posterior likelihood ratio satisfies

$$\frac{\mu_s(\{n-1,n+1\})}{\mu_s(n)} = \frac{f(s|n-1)}{f(s|n)}\frac{\mu(n-1)}{\mu(n)} + \frac{f(s|n+1)}{f(s|n)}\frac{\mu(n+1)}{\mu(n)} > c\left(\frac{f(s|n-1)}{f(s|n)} + \frac{f(s|n+1)}{f(s|n)}\right).$$

As the last expression is the sum of a strictly positive decreasing function of s and a strictly positive increasing function of s, it is bounded away from 0 in s. The bound is independent of n because normal information is a location-shift family of distributions. Therefore, the posterior likelihood ratio is uniformly bounded away from zero, and hence, the posterior $\mu_s(n)$ is uniformly bounded away from one.

[24] Since normal information has standard density $g(s) = \frac{1}{\sigma\sqrt{2\pi}}\exp[-(1/2)(s/\sigma)^2]$ for some $\sigma > 0$, it need not hold that $g(s) \leq \exp[-|s|^2]$ for all large s in absolute value, but for any $p \in (1,2)$, $g(s) \leq \exp[-|s|^p]$ for all large s in absolute value.

dependent of s.[25] More broadly, thick-tailed standard densities will tend to violate DUB, intuitively because a thick tail implies that fixing any pair of states, extreme signals become uninformative.[26] For example, DUB fails for standardized student t-distribution with arbitrary degrees of freedom: it can be checked that for any $\omega' < \omega$, as $s \to \infty$ or $s \to -\infty$, $f(s|\omega')/f(s|\omega) \to 1$. By Proposition 1 , such signal structures will entail inadequate learning under some SCD preferences.

DUB given MLRP. Another common assumption on signal structures is that of MLRP (fn. 11). Say that there is *pairwise unbounded beliefs* if any two states are distinguishable from each other (Arieli and Mueller-Frank, 2021). Given MLRP, it is straightforward to check for DUB because:

Remark 1. Given MLRP, DUB is equivalent to pairwise unbounded beliefs.

Plainly, regardless of MLRP, DUB implies pairwise unbounded beliefs. For the intuition why the converse is true given MLRP, consider the case of finite states. Note that for any $\omega' > \omega$, $f(s|\omega')/f(s|\omega) \to \infty$ as $s \to \sup S$ (the ratio is increasing by MLRP, and it diverges by pairwise unbounded beliefs); similarly, the ratio goes to 0 as $s \to \inf S$. Hence, for any ω and $\varepsilon > 0$, condition 1 of Lemma 1 is met when $\overline{S}$ is any sufficiently small upper set of signals, and condition 2 is met when $\underline{S}$ is any sufficiently small lower set. For an infinite state space, the intuition is the same but we must appeal to the monotone convergence theorem; see fn. 35.

5.2. Learning at a Fixed Prior

Given SCD preferences, Proposition 1 says that DUB is necessary when learning is required for all priors. One may also be interested in whether DUB is necessary for learning at some given prior, in particular at an arbitrary full-support prior. When there are only two states, there is no distinction between these two concepts: if learning fails at any prior, then there is an open ball of stationary beliefs around certainty on some state; hence, no matter the (full-support) prior, there cannot be a belief path that converges to certainty on that state.

[25] In this example g is logconcave, which is equivalent to the location-shift signal structure having the MLRP. So any state ω is distinguishable from a lower state ω' if and only if $f(s|\omega')/f(s|\omega) = g(s-\omega')/g(s-\omega) \to 0$ as $s \to \infty$.

[26] More precisely, if g is superexponential in the sense that there is some $p < 1$ such that for all s large enough, $g(s) \geq \exp(-s^p)$, then one can show that for any $z > 0$, $\lim_{s\to\infty} g(s+z)/g(s) = 1$ if the limit exists. That is, for any two states that differ by z, extremely high signals do not distinguish them. The argument is similar to that used in proving Proposition 3.

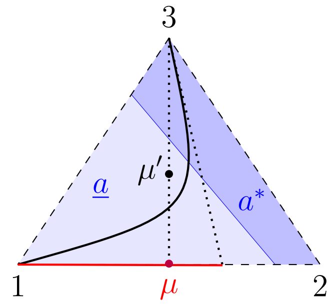

Figure 3: Preference regions among the actions $\underline{a}$ and a^* shaded in blues. An MLRP signal structure such that under belief μ', posteriors are given by the black curve, while under belief μ, posteriors are given by the red line.

However, with multiple states, a failure of learning at some prior does not not imply an open ball of stationary beliefs around certainty on any state, even under SCD preferences. To illustrate, consider Figure 3. Action a^* is optimal at states 2 and 3 while $\underline{a}$ is optimal at state 1. Learning fails when the prior is μ because μ, which has support $\{1, 2\}$, is stationary but has inadequate knowledge. (SCD holds, but DUB is violated because state 2 is not distinguishable from state 1 or state 3.) Yet there is no open ball of stationary beliefs around certainty on any state. Indeed, no full-support belief is stationary because the optimal actions are distinct at the extreme states 1 and 3, and the extreme states are distinguishable from their complements. This raises the possibility that there is learning for some—or even all—full-support priors, with on-path sequences of beliefs (which necessarily have full support at every finite time) converging to certainty without ever hitting any stationary belief (all of which have non-full-support).

Nevertheless, we can establish that, at least under some conditions, a failure of DUB implies a failure of learning for some SCD preference at every full-support prior:

Proposition 4. *Suppose the network is complete. If the signal structure has MLRP but fails DUB, then there is an SCD preference such that there is inadequate learning for all full-support priors.*

In particular, under a complete network, the signal structure and preferences in Figure 3 entail inadequate learning at every full-support prior, such as μ' in the figure.[27]

Here is the key idea behind Proposition 4. For simplicity, suppose $\Omega = \mathbb{Z}$. We will argue that if for every SCD preference there is a full-support prior for which there is adequate learning, then DUB holds. Pick any state ω^*, and consider SCD preferences such that a' is strictly preferred to a^* at every state $\omega < \omega^*$ while a^* is strictly preferred at every $\omega \geq \omega^*$

[27] It can be confirmed that the figure's signal structure satisfies MLRP because the black curve is concave vis-à-vis the 1–3 edge and approaches the 1 and 3 vertices.

(and all other actions are dominated, hence can be ignored). Assume there is adequate learning at some full-support prior. It can be shown that MLRP and the complete network then guarantee the following property: with positive probability there is an infinite history with a herd on a', call it $h' \equiv (\underbrace{a_1, \ldots, a_n}_{\text{each } a_i \in \{a', a^*\}}, a', a', \ldots)$, whose limit belief μ' satisfies $\mathrm{Supp}\, \mu' = \{\omega : \omega < \omega^*\}$. This support property can be restated as $\Pr(h'|\omega) > 0 \iff \omega < \omega^*$. So the probability of signals that overturn the herd—i.e., lead to action a^*—must vanish over time at a fast enough rate in each $\omega < \omega^*$, but either not vanish or vanish at a slow enough rate in each $\omega \geq \omega^*$. In particular, there must exist overturning signals whose probability gets arbitrarily large in state ω^* relative to those in every $\omega < \omega^*$, which implies that ω^* is distinguishable from its lower set $\{\omega : \omega < \omega^*\}$. A symmetric argument using preferences in which a^* is optimal in state ω^* and below while a' is optimal in states above ω^* establishes that ω^* is distinguishable from its upper set $\{\omega : \omega > \omega^*\}$. Since ω^* was arbitrary, it follows that there is DUB.

While Proposition 4 is stated for the complete network and MLRP signal structures, we suspect that the argument is more general. These conditions are only used in combination with learning to obtain a herd on a' that has positive probability in each state in which this action is optimal. Specifically, the complete network guarantees that there is almost surely a herd on a' in state $\omega^* - 1$, and MLRP ensures that any history with a herd on a' that has positive probability in state $\omega^* - 1$ also has positive probability in all lower states.[28] We hope future research will settle to what extent these properties hold more generally.

5.3. Excludability

In explaining Theorem 1, we highlighted that what is crucial for learning is that a single agent should always be able to rule out a wrong action, rather than take the correct action. This subsection expands on a general condition on preferences and information—moving beyond SCD and DUB—that guarantees this "ruling out" property. Beyond its conceptual value, we believe this will be useful for future research investigating other families of

[28] The complete network assumption implies that the public belief—an agent's belief based on the history before she observes her private signal–is a martingale, and hence converges almost surely. If the limit public belief has adequate knowledge, then either the public belief on $\{\omega : \omega < \omega^*\}$ goes to 0, in which case there is a herd on a', or the public belief on $\{\omega : \omega \geq \omega^*\}$ goes to 0, in which case there is a herd on a^*. Since there is adequate learning, there must be a herd almost surely; this is a manifestation of Smith and Sørensen's (2000) overturning principle. Adequate learning implies that in any state $\omega < \omega^*$ the herd must be on a', while in any state $\omega \geq \omega^*$ the herd must be on a^*. Hence, following any herd on a', the limit public belief excludes all $\omega \geq \omega^*$. Take any infinite history with a herd on a' that has positive probability in state $\omega^* - 1$ (such a history exists). MLRP implies that a' is only chosen for a lower set of signals and that the probability of any lower set of signals is nonincreasing in the state. Thus, the history has positive probability under every $\omega < \omega^*$, and so the limit belief's support is $\{\omega : \omega < \omega^*\}$.

preferences and information that generate learning.

For an arbitrary utility function $u(a,\omega)$ and any two actions a_1 and a_2, define the dominance set $\Omega_{a_1,a_2} := \{\omega : \; u(a_1,\omega) - u(a_2,\omega) > 0\}$ as the set of states in which a_1 is strictly preferred to a_2. Say that a utility function and information structure jointly satisfy *excludability* if for every pair of actions a_1 and a_2, every $\omega \in \Omega_{a_1,a_2}$ is distinguishable from Ω_{a_2,a_1}. That is, excludability holds when, for any pair of actions, it is possible to become arbitrarily certain that one action is strictly better than the other starting from any belief that does not exclude that event.

Proposition 5. *There is adequate learning if there is excludability. If excludability fails and there are a finite number of states, then there is inadequate learning for some choice set.*

Proposition 5 derives from Lemma 2. Although excludability is a joint condition on preferences and information, it is more interpretable and tractable than the lemma's condition that all stationary beliefs have adequate knowledge. For, Proposition 5 allows one to focus on whether information affords distinguishability of all pairs of dominance sets. To appreciate why this is useful, observe that an SCD preference is characterized by every dominance set being either an upper set or a lower set. Thus, an information structure yields excludability for all SCD preferences if and only if that information structure satisfies DUB.

Various other learning results in literature also amount to identifying conditions that imply excludability. For instance, an information structure yields excludability for all utility functions if and only if that information structure has unbounded beliefs. Another example is Arieli and Mueller-Frank (2021, Theorem 3), which establishes that pairwise distinguishability—for any pair of states, each is distinguishable from the other—is sufficient for learning if the agents have a simple utility function that gives utility 1 if the action matches the state and 0 otherwise. Here the dominance sets are all individual states, and so excludability reduces to the information structure having pairwise distinguishability.

The second portion of Proposition 5 says that if we can vary the choice set (and subject to the finiteness qualifier), then excludability is also necessary for learning. Varying the choice set is often subsumed when considering a family of preferences; for example, any SCD utility function restricted to a subset of actions still has SCD, so the class of all SCD preferences on the action set A automatically includes all choice sets $\tilde{A} \subset A$. Hence, Proposition 1 is a direct corollary of Proposition 5, up to the finiteness qualifier in the latter.[29]

[29] The appendix provides a stronger version of Proposition 5, Proposition 5′, that does not require finiteness and can be used to prove Proposition 1.

The intuition for the necessity of excludability is also similar to that of Proposition 1: if excludability fails, then there is some dominance set Ω_{a_1,a_2} that cannot be distinguished from Ω_{a_2,a_1}. This means that under the choice set $\tilde{A} = \{a_1, a_2\}$, a prior that puts small enough probability on Ω_{a_1,a_2} relative to Ω_{a_2,a_1} will be stationary yet have inadequate knowledge.

For a fixed choice set, excludability is not necessary for adequate learning. A well-known example is that of responsive preferences (Lee, 1993; Ali, 2018): there is a different optimal action for every belief about the state. For example, consider $\Omega = \{0, 1\}$, $A = [0, 1]$, and $u(a, \omega) = -(a - \omega)^2$. Adequate learning obtains in this example for any nontrivial information structure by Lemma 2, because given any nondegenerate belief, with positive probability the posterior-optimal action will be different from the prior-optimal action. However, excludability is equivalent to unbounded beliefs.

References

ACEMOGLU, D., M. A. DAHLEH, I. LOBEL, AND A. OZDAGLAR (2011): "Bayesian Learning in Social Networks," *Review of Economic Studies*, 78, 1201–1236.

AGHION, P., P. BOLTON, C. HARRIS, AND B. JULLIEN (1991): "Optimal Learning by Experimentation," *Review of Economic Studies*, 58, 621–654.

ALI, S. N. (2018): "On the Role of Responsiveness in Rational Herds," *Economics Letters*, 163, 79–82.

ARIELI, I. AND M. MUELLER-FRANK (2021): "A General Analysis of Sequential Social Learning," *Mathematics of Operations Research*, 46, 1235–1249.

ATHEY, S. (2001): "Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information," *Econometrica*, 69, 861–889.

——— (2002): "Monotone Comparative Statics under Uncertainty," *Quarterly Journal of Economics*, 117, 187–223.

AUMANN, R. J. AND M. MASCHLER (1995): *Repeated Games with Incomplete Information*, MIT Press.

AVERY, C. AND P. ZEMSKY (1998): "Multi-Dimensional Uncertainty and Herd Behavior in Financial Markets," *American Economic Review*, 88, 724–748.

BANERJEE, A. AND D. FUDENBERG (2004): "Word-of-mouth Learning," *Games and Economic Behavior*, 46, 1–22.

BANERJEE, A. V. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107, 797–817.

BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100, 992–1026.

BOGACHEV, V. I. (2018): *Weak Convergence of Measures*, American Mathematical Society Providence.

BOHREN, J. A. AND D. HAUSER (2021): "Learning with Model Misspecification: Characterization and Robustness," Forthcoming in *Econometrica*.

ÇELEN, B. AND S. KARIV (2004): "Observational Learning under Imperfect Information," *Games and Economic Behavior*, 47, 72–86.

CRAUEL, H. (2002): *Random Probability Measures on Polish Spaces*, vol. 11, CRC press.

DURRETT, R. (2019): *Probability: Theory and Examples*, vol. 49, Cambridge University Press.

EYSTER, E., A. GALEOTTI, N. KARTIK, AND M. RABIN (2014): "Congested Observational Learning," *Games and Economic Behavior*, 87, 519–538.

FRICK, M., R. IIJIMA, AND Y. ISHII (2020): "Misinterpreting Others and the Fragility of Social Learning," *Econometrica*, 88, 2281–2328.

——— (2021): "Belief Convergence under Misspecified Learning: A Martingale Approach," Unpublished.

GALE, D. AND S. KARIV (2003): "Bayesian Learning in Social Networks," *Games and Economic Behavior*, 45, 329–346.

GOEREE, J. K., T. R. PALFREY, AND B. W. ROGERS (2006): "Social Learning with Private and Common Values," *Economic Theory*, 26, 245–264.

GOLUB, B. AND E. SADLER (2016): "Learning in Social Networks," in *The Oxford Handbook of the Economics of Network*, ed. by Y. Bramoullé, A. Galeotti, and B. W. Rogers, Oxford and New York: Oxford University Press, 504–542.

HANN-CARUTHERS, W., V. V. MARTYNOV, AND O. TAMUZ (2018): "The Speed of Sequential Asymptotic Learning," *Journal of Economic Theory*, 173, 383–409.

JAFFRAY, J.-Y. (1975): "Semicontinuous Extension of a Partial Order," *Journal of Mathematical Economics*, 2, 395–406.

KAMENICA, E. AND M. GENTZKOW (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

KARTIK, N., S. LEE, AND D. RAPPOPORT (2019): "Single-Crossing Differences on Distributions," Unpublished.

LEE, I. H. (1993): "On the Convergence of Informational Cascades," *Journal of Economic Theory*, 61, 395–411.

LOBEL, I. AND E. SADLER (2015): "Information Diffusion in Networks through Social Learning," *Theoretical Economics*, 10, 807–851.

MILGROM, P. AND C. SHANNON (1994): "Monotone Comparative Statics," *Econometrica*, 62, 157–180.

MILGROM, P. R. (1979): "A Convergence Theorem for Competitive Bidding with Differential Information," *Econometrica*, 47, 679–688.

QUAH, J. K.-H. AND B. STRULOVICI (2009): "Comparative Statics, Informativeness, and the Interval Dominance Order," *Econometrica*, 77, 1949–1992.

ROSENBERG, D. AND N. VIEILLE (2019): "On the Efficiency of Social Learning," *Econometrica*, 87, 2141–2168.

SMITH, L. AND P. SØRENSEN (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68, 371–398.

——— (2020): "Rational Social Learning with Random Sampling," Unpublished.

A. Backbone Results

We prove our backbone results—Theorem 2, Lemma 2, and Proposition 5—in the following setting, which is more general than that described in the main text. These results do not require SCD or DUB, nor any order structure on Ω.

- The action space and signal space $(A, \mathcal{A}), (S, \mathcal{S})$ are standard Borel spaces;
- The state space Ω is equipped with a metric d and its Borel sigma-algebra, $\mathcal{B}(\Omega)$, such that (Ω, d) is a sigma-compact Polish space;[30,31]
- The utility function $u(a, \omega)$ is uniformly bounded by $\overline{u}$ and is pointwise equicontinuous when regarded as a collection of functions of ω indexed by a; moreover, for every belief (Borel probability distribution over Ω), there exists an optimal action;
- The information/signal structure $F(\cdot|\omega)$ is a Markov kernel from $(\Omega, \mathcal{B}(\Omega))$ to $(S, \mathcal{S})$ that is continuous in ω in the total variation (TV) sense;
- The network structure $Q \equiv (Q_n)_{n\in\mathbb{N}}$, where each Q_n is a probability measure over all neighborhoods, i.e., all subsets of $\{1, 2, \ldots, n-1\}$, independent across n, independent of the state ω, and independent of any private signals.

Note that the main text's setting is subsumed because our sigma-compactness and continuity requirements are trivially satisfied when a countable Ω is endowed with the discrete topology.

With these elements, we can define an overarching probability space over all realizations of the state, signals, observation neighborhoods, and actions. An overarching probability measure is determined uniquely by the prior μ_0, the signal structure F, the network structure Q, and a strategy profile σ. We will refer to this overarching probability measure by $\mathbb{P}_{\sigma,\mu_0}$, or sometimes just $\mathbb{P}$ for short. The agents' beliefs are defined as the regular conditional probabilities of this overarching probability measure, which are guaranteed to exist and are unique almost everywhere. Appendix D elaborates.

We will denote by $\Delta\Omega$ the space of beliefs (Borel probability measures on Ω) equipped with the Prohorov metric, and by $\Delta\Delta\Omega$ the space of belief distributions (Borel probability measures on $\Delta\Omega$) equipped with the Prohorov metric.

[30] That is, (Ω, d) is a complete and separable metric space that can be represented as a countable union of compact sets.

[31] Our backbone results hold regardless of the topology on Ω, so long as the sigma-compactness and subsequent continuity assumptions are met. The topology only plays a role in the proofs. Indeed, given some utility function and information structure, the backbone results hold so long as we can find a topology ensuring the compactness/continuity properties.

A.1. Space of Bayes-plausible Belief Distributions is Compact

Given a prior $\mu_0 \in \Delta\Omega$ and a strategy profile σ, any agent's belief distribution $\varphi \in \Delta\Delta\Omega$ must be Bayes plausible:

$$\int_A \mu(A)\,\mathrm{d}\varphi(\mu) = \mu_0(A),\ \forall A \in \mathcal{B}(\Omega). \tag{1}$$

Let $\Phi^{BP} \subset \Delta\Delta\Omega$ be the set of Bayes-plausible belief distributions; note that we suppress the dependence of Φ^{BP} on μ_0.

Our goal is to establish (Lemma 4 below) that even though the set of belief distributions $\Delta\Delta\Omega$ need not be compact, the subset of Bayes-plausible distributions Φ^{BP} is. A key step is the following lemma, which shows that any belief $\varphi \in \Phi^{BP}$ has to put a large probability on a compact subset of $\Delta\Omega$.

Lemma 3. *Let $\delta > 0$ and $\{\Omega_i\}_{i\in\mathbb{N}}$ be a sequence of compact sets with $\mu_0(\Omega_i) \geq 1 - (\frac{\delta}{2^i})^2$, $\forall i$. Defining $V_\delta := \{\mu \in \Delta\Omega : \mu(\Omega_i) \geq 1 - \frac{\delta}{2^i},\ \forall i\}$, it holds that:*

1. *V_δ is compact;*
2. *$\varphi(\mu \notin V_\delta) < \delta, \forall \varphi \in \Phi^{BP}$.*

Intuitively, in the lemma's statement, the set V_δ contains all beliefs that put high probability on a set of states that the prior μ_0 ascribes high probability to. The lemma concludes that the set V_δ is compact and that any Bayes-plausible belief distribution must put high probability on V_δ.

Proof. (Part 1) First, V_δ is closed. To see this, take any $\mu_k \to \mu$ and $\mu_k \in V_\delta$. Since each Ω_i is compact (and thus closed), weak convergence implies

$$\limsup_k \mu_k(\Omega_i) \leq \mu(\Omega_i),\ \forall i,$$

which implies $\mu(\Omega_i) \geq 1 - \frac{\delta}{2^i}$. Thus, $\mu \in V_\delta$, and hence V_δ is closed.

Next, the beliefs in V_δ are tight by definition. Hence, by Prohorov's Theorem, the closure of V_δ, which is V_δ itself, is compact.

(Part 2) Note that $\varphi(\mu \notin V_\delta) = \varphi(\cup_i\{\mu(\Omega_i^c) > \frac{\delta}{2^i}\}) \leq \sum_i \varphi(\mu(\Omega_i^c) > \frac{\delta}{2^i})$. For each $i \in \mathbb{N}$, we view $\mu(\Omega_i^c)$ as a non-negative random variable with a distribution φ. Since φ is Bayes plausible, $\mathbb{E}_\varphi[\mu(\Omega_i^c)] = \mu_0(\Omega_i^c) \leq (\frac{\delta}{2^i})^2$, which implies (using Markov's inequality) that $\varphi(\mu(\Omega_i^c) > \frac{\delta}{2^i}) < \frac{\delta}{2^i}$. This implies that $\varphi(\mu \notin V_\delta) < \sum_i \frac{\delta}{2^i} = \delta$. *Q.E.D.*

Given Lemma 3, the idea now is to apply Prohorov's theorem again to Φ^{BP} and obtain:

Lemma 4. *Φ^{BP} is compact.*

Proof. First, to prove that Φ^{BP} is closed, it is sufficient to take any $\varphi_k \to \varphi$ and $\varphi_k \in \Phi^{BP}$, and show that $\varphi \in \Phi^{BP}$, i.e., $\mathbb{E}_\varphi[\mu(A)] = \mu_0(A), \forall A \in \mathcal{B}(\Omega)$.

Take any open set $A \in \mathcal{B}(\Omega)$. For any $\mu_k \to \mu$, it holds that $\mu(A) \leq \liminf \mu_k(A)$. In other words, $\mu(A)$ (as a function of μ) is lower semi-continuous. By properties of weak convergence, it follows that $\mathbb{E}_\varphi[\mu(A)] \leq \liminf \mathbb{E}_{\varphi_k}[\mu(A)] = \mu_0(A)$. That is, the mean measure of φ ascribes a smaller probability than μ_0 to any open set.

Now observe that $A^c \subset \cup_{x\in A^c} B_{1/n}(x)$ for any n. Hence,

$$\mathbb{E}_\varphi[\mu(A^c)] \leq \lim_n \mathbb{E}_\varphi[\mu(\cup_{x\in A^c} B_{1/n}(x))] \leq \lim_n \mu_0(\cup_{x\in A^c} B_{1/n}(x)) = \mu_0(A^c),$$

where the second inequality is from the previous result applied to open sets $\cup_{x\in A^c} B_{1/n}(x)$, and the last equality follows from $A^c = \cap_n \cup_{x\in A^c} B_{1/n}(x)$ (and this equality holds because A^c is closed). Therefore, $\mathbb{E}_\varphi[\mu(A)] = \mu_0(A)$.

Since $\mathbb{E}_\varphi[\mu]$ and μ_0 agree on all open sets, and open sets generate $\mathcal{B}(\Omega)$, $E_\varphi[\mu]$ and μ_0 agree on all sets in $\mathcal{B}(\Omega)$.

Finally, Ω being sigma-compact implies that for any δ, there is an increasing sequence of compact sets $\{\Omega_i\}_{i\in\mathbb{N}}$ such that $\Omega = \cup_i \Omega_i$, and so this sequence $\{\Omega_i\}$ satisfies the hypotheses in Lemma 3. The lemma guarantees that there is a compact set V_δ such that $\varphi(V_\delta) < \delta$ for all $\varphi \in \Phi^{BP}$, and so Φ^{BP} is tight. Prohorov's theorem now implies that $\mathrm{cl}(\Phi^{BP}) = \Phi^{BP}$ is compact. *Q.E.D.*

A.2. Continuity of Various Functions

We next define some functions of interest, some of which were already defined in the main text but are now defined for the more general setting considered in the appendix.

Let $u(\mu)$ be the expected utility that an agent can get at belief μ:

$$u(\mu) := \sup_a \int_\Omega u(a, \omega)\,\mathrm{d}\mu(\omega).$$

Let $u^F(\mu)$ be the expected utility that an agent can get at belief μ, if she can choose an action after observing her private signal:

$$u^F(\mu) := \sup_{\beta:S\to A} \int_\Omega \int_S u(\beta(s), \omega)\,\mathrm{d}F(s|\omega)\,\mathrm{d}\mu(\omega).$$

Finally, let $u^*(\mu)$ be the full information utility at μ:

$$u^*(\mu) := \int_\Omega \sup_a u(a, \omega) \, \mathrm{d}\mu(\omega).$$

Our continuity assumptions on the utility function and the information structure allow us to prove:

Lemma 5. *u, u^F, u^* are continuous in μ.*

To prove Lemma 5, we require Theorem 2.2.8 in Bogachev (2018), which we restate without proof for our context as the following claim:

Claim 1. *Let $\mu_k \to \mu$. If Γ is a uniformly bounded and pointwise equicontinuous family of functions on Ω, then*

$$\lim_k \sup_{f \in \Gamma} \left| \int_\Omega f \, \mathrm{d}\mu_k - \int_\Omega f \, \mathrm{d}\mu \right| = 0.$$

Proof of Lemma 5. By assumption, $\Gamma := \{u(a, \omega)\}_{a \in A}$, viewed as a family of functions of ω indexed by a, is uniformly bounded and pointwise equicontinuous.

Consider the function u^*. Since the supremum of the pointwise equicontinuous functions $u^*(\omega) := \sup_a u(a, \omega)$ is continuous in ω, the definition of weak convergence implies that $u^*(\mu)$ is continuous in μ.

Now consider the function u. Its continuity follows from

$$|u(\mu_k) - u(\mu)| = \left| \sup_{f \in \Gamma} \int_\Omega f \, \mathrm{d}\mu_k - \sup_{f \in \Gamma} \int_\Omega f \, \mathrm{d}\mu \right| \leq \sup_{f \in \Gamma} \left| \int_\Omega f \, \mathrm{d}\mu_k - \int_\Omega f \, \mathrm{d}\mu \right|,$$

which converges to 0 for $\mu_k \to \mu$ by Claim 1.

Lastly, suppose we establish that $\Gamma^F := \left\{ \int_S u(\beta(s), \omega) \, \mathrm{d}F(s|\omega) \right\}_{\beta: S \to A}$, as a family of functions of ω indexed by β, is pointwise equicontinuous.[32] Then, as Γ^F is uniformly bounded, Claim 1 implies that $u^F(\mu)$ is continuous, proving the lemma.

To establish the pointwise equicontinuity of Γ^F, observe that $\forall \omega, \omega'$ and $\forall \beta$,

$$\begin{aligned} &\left| \int_S u(\beta(s), \omega) \, \mathrm{d}F(s|\omega) - \int_S u(\beta(s), \omega') \, \mathrm{d}F(s|\omega') \right| \qquad (2) \\ \leq &\left| \int_S (u(\beta(s), \omega) - u(\beta(s), \omega')) \, \mathrm{d}F(s|\omega) \right| + \left| \int_S u(\beta(s), \omega') \, \mathrm{d}F(s|\omega) - \int_S u(\beta(s), \omega') \, \mathrm{d}F(s|\omega') \right|. \end{aligned}$$

[32] Here we assume β are (measurable) pure strategies for notation clarity. The same argument works for mixed strategies, in which case β would be Markov kernels.

Fix any ω and any $\varepsilon > 0$. Since $\{u(a,\omega)\}_{a\in A}$ is pointwise equicontinuous, there exists δ_1 such that $d(\omega',\omega) < \delta_1$ implies the first term on the right-hand side of inequality (2) to be smaller than $\varepsilon/2$ (regardless of $\beta(s)$). The second term is smaller than $2\overline{u}d_{TV}(F(\cdot|\omega), F(\cdot|\omega'))$ (where TV represents total variation), and by the continuity assumption of the information structure, there exists $\delta_2 > 0$ such that $d(\omega',\omega) < \delta_2$ implies $d_{TV}(F(\cdot|\omega), F(\cdot|\omega')) < \varepsilon/4\overline{u}$. Therefore, if $d(\omega',\omega) < \min\{\delta_1, \delta_2\}$, then the right-hand side of inequality (2) is less than ε (regardless of $\beta(s)$). It follows that Γ^F is pointwise equicontinuous. *Q.E.D.*

Now define the utility improvement $I(\mu)$ and the utility gap $G(\mu)$ at μ as:

$$I(\mu) := u^F(\mu) - u(\mu), \quad G(\mu) := u^*(\mu) - u(\mu).$$

By Lemma 5, $I(\mu)$ and $G(\mu)$ are continuous. Lastly, with an abuse of notation, define $u(\varphi) := \mathbb{E}_\varphi[u(\mu)]$, $I(\varphi) := \mathbb{E}_\varphi[I(\mu)]$, and $G(\varphi) := \mathbb{E}_\varphi[G(\mu)]$ as the corresponding functions over distributions of beliefs. Since $u(\mu)$, $I(\mu)$, and $G(\mu)$ are continuous, so are $u(\varphi)$, $I(\varphi)$, and $G(\varphi)$.

A.3. Proofs of Backbone Results

We say that a belief μ is *stationary* if $I(\mu) = 0$, and a belief μ has *adequate knowledge* if $G(\mu) = 0$. Note that these definitions agree with those in the main text. To see that, consider stationary beliefs. If there is an action that is a.s. optimal regardless of the signal, then clearly $I(\mu) = 0$. Conversely, if there is no action that is a.s. optimal regardless of the signal, then for any action there is a positive-probability set of signals for which that action is strictly suboptimal; hence $u^F(\mu) > u(\mu)$, and $I(\mu) > 0$. The argument for adequate knowledge beliefs is similar and omitted.

Proof of Theorem 2. We prove the result in two steps. In Step 1 below, we prove that if agent n's belief distribution φ_n, which is her belief distribution incorporating the observation of her neighborhood's actions but not her private signal, is not close to being supported on only stationary beliefs, then her utility $\mathbb{E}_{\sigma,\mu_0}[u_n]$, which is the ex-ante expected utility under equilibrium σ after observing the private signal, improves from $u(\varphi_n)$ by a fixed positive minimum. In Step 2 below, we use the expanding observations assumption to establish that this minimum improvement propagates through the network until eventually agents obtain at least arbitrarily close to their cascade utility level.

Step 1: Recall the set of Bayes-plausible belief distributions that are supported by stationary beliefs, $\Phi^S := \{\varphi \in \Phi^{BP} : I(\varphi) = 0\}$, and the *cascade utility*, $u_* := \inf_{\varphi\in\Phi^S} u(\varphi)$.

Take any $\varepsilon > 0$, and let $(\Phi^S)^\varepsilon$ denote the ε-neighborhood of Φ^S. An agent n's belief distribution φ_n must be Bayes plausible, so $\varphi_n \in \Phi^{BP}$. Since $u(\varphi)$ is uniformly continuous (being continuous on the compact set Φ^{BP}), if $\varphi_n \in (\Phi^S)^\varepsilon$, then $u(\varphi_n) \geq u_* - \gamma(\varepsilon)$ for some $\gamma(\cdot)$ such that $\gamma(\varepsilon) \to 0$ when $\varepsilon \to 0$. If, on the other hand, $\varphi_n \in \Phi^+ := \Phi^{BP} \backslash (\Phi^S)^\varepsilon$, then $I(\varphi_n) > 0$ because φ_n puts positive probability on $\{\mu : I(\mu) > 0\}$. Note that Φ^+ is a closed subset of a compact set Φ^{BP}, so it is compact, and since $I(\varphi)$ is continuous, it attains a minimum over Φ^+ at some $\underline{\varphi} \in \Phi^+$. Thus, if $\varphi_n \in \Phi^+$ the agent obtains a minimum improvement $I(\varphi_n) \geq I(\underline{\varphi}) > 0$.

Step 2: We will argue that for any $\varepsilon > 0$, $\mathbb{E}_{\sigma,\mu_0}[u_n] \geq u_* - \gamma(\varepsilon)$ once n is large enough. Since ε is arbitrary, taking $\varepsilon \to 0$ implies $\liminf \mathbb{E}_{\sigma,\mu_0}[u_n] \geq u_*$, which completes the proof.

For a given $\varepsilon > 0$, let $\delta = \frac{I(\underline{\varphi})}{4\overline{u}} > 0$, let $N_0 = 1$, and define N_k for $k = 1, 2, \ldots$ sequentially such that for all $n \geq N_k$, $Q_n(\max_{b \in B_n} b < N_{k-1}) < \delta$. Expanding observations ensures that such N_k exist.

We claim that, for any agent $n \geq N_k$, $\mathbb{E}_{\sigma,\mu_0}[u_n] \geq \alpha_k := \min\{u_* - \gamma(\varepsilon), \frac{kI(\underline{\varphi})}{2} - \overline{u}\}$. Since $\alpha_0 = -\overline{u}$, clearly $\mathbb{E}_{\sigma,\mu_0}[u_n] \geq \alpha_0$ for any $n \geq N_0$. Suppose the claim holds for all agents $n' \geq N_{k-1}$. Take any agent $n \geq N_k$. Agent n's neighborhood is drawn independently of everything that has happened before, so conditional on agent n observing an agent $n' \geq N_{k-1}$, even without her private signal agent n can achieve a utility of at least α_{k-1} by imitating agent n'. Hence, $u(\varphi_n) \geq (1-\delta) \cdot \alpha_{k-1} + \delta \cdot (-\overline{u})$. If $\varphi_n \in (\Phi^S)^\varepsilon$, then by definition $u(\varphi_n) \geq u_* - \gamma(\varepsilon)$, and thus $\mathbb{E}_{\sigma,\mu_0}[u_n] \geq u(\varphi_n) \geq u_* - \gamma(\varepsilon) \geq \alpha_k$. If $\varphi_n \notin (\Phi^S)^\varepsilon$, then agent n can improve her utility by at least $I(\underline{\varphi})$, and so

$$
\begin{aligned}
\mathbb{E}_{\sigma,\mu_0}[u_n] &\geq (1-\delta)\alpha_{k-1} + \delta \cdot (-\overline{u}) + I(\underline{\varphi}) \\
&\geq \alpha_{k-1} + \frac{I(\underline{\varphi})}{2} \quad \text{(because } \alpha_{k-1} \leq \overline{u} \text{ and } \delta = \frac{I(\underline{\varphi})}{4\overline{u}}\text{)} \\
&\geq \alpha_k.
\end{aligned}
$$

Since the definition of α_k implies that there is a finite K such that for all $k \geq K$, $\alpha_k = u_* - \gamma(\varepsilon)$, it follows that for all $n \geq N_K$, $\mathbb{E}_{\sigma,\mu_0}[u_n] \geq u_* - \gamma(\varepsilon)$. *Q.E.D.*

Proof of Lemma 2. The "only if" direction is straightforward. If there is a stationary belief without adequate knowledge, then when the prior is that belief there is an equilibrium where each agent ignores her signal and action history and obtains a utility that is strictly below the full-information utility level.

For the "if" direction, fix any prior μ_0 and equilibrium σ. Since all stationary beliefs have adequate knowledge, $I(\mu) = 0$ implies $G(\mu) = 0$. Thus, for any $\varphi \in \Phi^S$, $\varphi(\{\mu : I(\mu) = 0\}) =$

$\varphi(\{\mu : G(\mu) = 0\}) = 1$, which implies $G(\varphi) = u^*(\varphi) - u(\varphi) = 0$. Moreover, because μ_0 is the mean measure of φ,

$$u^*(\varphi) = \mathbb{E}_\varphi \left[\int_\Omega \sup_a u(a,\omega)\, \mathrm{d}\mu \right] = \int_\Omega \sup_a u(a,\omega)\, \mathrm{d}\mu_0 = u^*(\mu_0),$$

which implies $u(\varphi) = u^*(\mu_0)$. As a result, $u_*(\mu_0) = \inf_{\varphi \in \Phi^S} u(\varphi) = u^*(\mu_0)$. It follows from Theorem 2 that $\liminf_n \mathbb{E}_{\sigma,\mu_0}[u_n] \geq u^*(\mu_0)$. Since $\mathbb{E}_{\sigma,\mu_0}[u_n] \leq u^*(\mu_0)$ for all n, it further follows that $\mathbb{E}_{\sigma,\mu_0}[u_n] \to u^*(\mu_0)$. As μ_0 and σ are arbitrarily, we have adequate learning. *Q.E.D.*

Extending the notion given in the main text, for a general state space we say that a set Ω' is *distinguishable* from another set Ω'' if for any $\varepsilon > 0$ and any $\mu \in \Delta(\Omega' \cup \Omega'')$ with $\mu(\Omega') > 0$, there is a set of signals S' such that $\mathbb{P}_\mu(S') > 0$ and $\mu_s(\Omega') > 1 - \varepsilon$ for all $s \in S'$. Also generalizing from the main text, we say that preference u and signal structure F satisfy *excludability* if Ω_{a_1,a_2} is distinguishable from Ω_{a_2,a_1} for any $a_1, a_2 \in A$. Note that when the state space Ω is countable, as assumed the main text, Ω' is distinguishable from Ω'' if and only if every $\omega \in \Omega'$ is distinguishable from Ω''.

Next we state and prove a more general version of Proposition 5. For any $n \in \mathbb{N}$, define $\Omega^n_{a_1,a_2} := \{\omega : u(a_1,\omega) - u(a_2,\omega) > \frac{1}{n}\}$.

Proposition 5′. *There is adequate learning if*

$$\Omega^n_{a_1,a_2} \text{ is distinguishable from } \Omega_{a_2,a_1} \text{ for all } a_1, a_2 \text{ and } n. \tag{3}$$

There is inadequate learning for choice set $\{a_1, a_2\}$ *if*

$$\Omega_{a_1,a_2} \text{ is not distinguishable from } \Omega^n_{a_2,a_1} \text{ for some } n. \tag{4}$$

When Ω is countable, condition (3) is equivalent to excludability. When Ω is finite, or the utility difference between between any pair of actions is bounded away from zero, a failure of excludability is equivalent to condition (4) holding for some a_1, a_2.

Proof. (First statement) Lemma 2 implies that we need only show that any $\mu \in \Delta\Omega$ with inadequate knowledge is not stationary. So take any $\mu \in \Delta\Omega$ with inadequate knowledge and any $a \in c(\mu)$. Since there is inadequate knowledge, $\mu(\cup_{a' \in A} \Omega_{a',a}) > 0$, i.e., there is a positive measure of states where a is not optimal . The continuity of $u(a',\omega) - u(a,\omega)$ implies that $\Omega^n_{a',a}$ are open sets for any a' and n. Since Ω is Polish, it is second-countable and hence has a countable basis. Therefore, each open set $\Omega^n_{a',a}$, and hence the open set

$\cup_{a'}\Omega_{a',a}(=\cup_{a'}\cup_n\Omega^n_{a',a})$, is the union of countably many basic open sets. Since $\mu(\cup_{a'}\Omega_{a',a})>0$, at least one basic open set contained in $\Omega^n_{a',a}$ for some a' and n has strictly positive measure, i.e., $\mu(\Omega^n_{a',a})>0$.

Now denote $\mu'(\cdot):=\mu(\cdot|\Omega_{a,a'}\cup\Omega^n_{a',a})$ as the corresponding conditional probability. Since $\Omega^n_{a',a}$ is distinguishable from $\Omega_{a,a'}$ by assumption, for any $\varepsilon>0$ there exists a set of signals S' such that $\mathbb{P}_{\mu'}(S')>0$ and $\mu'_s(\Omega^n_{a',a})>1-\varepsilon$ for all $s\in S'$. The utility improvement upon observing any $s\in S'$ by switching from a to a' is therefore bounded below by $(\frac{1}{n}(1-\varepsilon)-2\overline{u}\varepsilon)\mu_s(\Omega_{a,a'}\cup\Omega^n_{a',a})$, as the expected improvement on $\Omega\backslash(\Omega_{a,a'}\cup\Omega^n_{a',a})$ is nonnegative. For small $\varepsilon>0$, $(\frac{1}{n}(1-\varepsilon)-2\overline{u}\varepsilon)>0$, and hence, integrating over S', the ex-ante improvement is bounded below by $(\frac{1}{n}(1-\varepsilon)-2\overline{u}\varepsilon)\mathbb{P}_{\mu'}(S')\mu(\Omega_{a,a'}\cup\Omega^n_{a',a})>0$. It follows that $I(\mu)>0$, and thus μ is not stationary.

(Second statement) Suppose there are two actions a_1,a_2 and an n such that Ω_{a_1,a_2} is not distinguishable from $\Omega^n_{a_2,a_1}$. This means there exists $\mu\in\Delta(\Omega_{a_1,a_2}\cup\Omega^n_{a_2,a_1})$ with $\mu(\Omega_{a_1,a_2})>0$ such that $\mu_s(\Omega_{a_1,a_2})\le 1-\varepsilon$ for some $\varepsilon>0$ and μ-a.e. s. Consider $\mu'\in\Delta(\Omega_{a_1,a_2}\cup\Omega^n_{a_2,a_1})$ with a small $\mu'(\Omega_{a_1,a_2})>0$ such that $\mu'(\cdot|\Omega_{a_1,a_2})=\mu(\cdot|\Omega_{a_1,a_2})$ and $\mu'(\cdot|\Omega^n_{a_2,a_1})=\mu(\cdot|\Omega^n_{a_2,a_1})$. Under μ', upon observing signal s, the posterior on Ω_{a_1,a_2} satisfies

$$\frac{\mu'_s(\Omega_{a_1,a_2})}{\mu'_s(\Omega^n_{a_2,a_1})}=\frac{\mu_s(\Omega_{a_1,a_2})/\mu(\Omega_{a_1,a_2})}{\mu_s(\Omega^n_{a_2,a_1})/\mu(\Omega^n_{a_2,a_1})}\frac{\mu'(\Omega_{a_1,a_2})}{\mu'(\Omega^n_{a_2,a_1})}\le\frac{1-\varepsilon}{\varepsilon}\frac{\mu(\Omega^n_{a_2,a_1})}{\mu(\Omega_{a_1,a_2})}\frac{\mu'(\Omega_{a_1,a_2})}{\mu'(\Omega^n_{a_2,a_1})}$$

for μ-a.e. s. Hence, by choosing μ' so that $\frac{\mu'(\Omega_{a_1,a_2})}{\mu'(\Omega^n_{a_2,a_1})}$ is arbitrarily small, the ratio $\frac{\mu'_s(\Omega_{a_1,a_2})}{\mu'_s(\Omega^n_{a_2,a_1})}$ can be made arbitrarily small uniformly over s.

Under μ', after observing s, the expected improvement by switching from a_2 to a_1 is bounded above by $2\overline{u}\mu'_s(\Omega_{a_1,a_2})-\frac{1}{n}\mu'_s(\Omega^n_{a_2,a_1})$, which is strictly negative when $\frac{\mu'_s(\Omega_{a_1,a_2})}{\mu'_s(\Omega^n_{a_2,a_1})}$ is small. Therefore, for μ'-a.e. s, a_2 is strictly better than a_1, and thus μ' is stationary for choice set $\{a_1,a_2\}$. However, since $\mu'(\Omega_{a_1,a_2})>0$, the belief μ' has inadequate knowledge. Lemma 2 implies there is inadequate learning for choice set $\{a_1,a_2\}$. *Q.E.D.*

B. Proofs for SCD & DUB being Minimally Sufficient

Given the general setup of Appendix A, we now assume in addition that Ω is totally ordered. To ensure that all intervals are open, we assume the topology on Ω is finer than the order topology. It is then without loss of generality to let $\Omega\subset\mathbb{R}$ because Ω can be embedded into $\mathbb{R}$ isomorphically.[33] It bears emphasizing that the topology on $\Omega\subset\mathbb{R}$ need not be the Euclidean subspace topology.

[33] For, it can be shown that Ω has a countable order-dense subset, and then Jaffray (1975, Corollary 1) applies. Details are available on request.

For sets $U, V \subseteq \Omega$, we say that U is smaller than a set V (denoted as $U < V$) if $u < v, \forall u \in U, v \in V$. Recalling the notion of distinguishability of two sets given before Proposition 5′, we define DUB for an arbitrary $\Omega \subset \mathbb{R}$ as follows:

Definition 2′. There is *directionally unbounded beliefs* (DUB) if all open sets $U, V \subset \Omega$ satisfying $U < V$ and $\mathrm{cl}(U) \cap \mathrm{cl}(V) = \emptyset$ are distinguishable from each other.

Note that this definition coincides with Definition 2 when Ω is countable and endowed with the discrete topology, but it applies more generally. We use this definition to prove the relevant results from the main text in the more general setting considered in this appendix.

Proof of Theorem 1. For any two actions a_1, a_2 and any $n \in \mathbb{N}$, both $\Omega^n_{a_1,a_2} \equiv \{\omega : u(a_1,\omega) - u(a_2,\omega) > \frac{1}{n}\}$ and Ω_{a_2,a_1} are open because of continuity, and $\mathrm{cl}(\Omega^n_{a_1,a_2}) \subset \{\omega : u(a_1,\omega) - u(a_2,\omega) \geq \frac{1}{n}\}$ and $\mathrm{cl}(\Omega_{a_2,a_1}) \subset \{\omega : u(a_2,\omega) - u(a_1,\omega) \geq 0\}$, which implies $\mathrm{cl}(\Omega^n_{a_1,a_2}) \cap \mathrm{cl}(\Omega_{a_2,a_1}) = \emptyset$. SCD implies $\Omega^n_{a_1,a_2} < \Omega_{a_2,a_1}$ or $\Omega^n_{a_1,a_2} > \Omega_{a_2,a_1}$. Either way, DUB implies $\Omega^n_{a_1,a_2}$ is distinguishable from Ω_{a_2,a_1}. It follows from Proposition 5′ that there is adequate learning. *Q.E.D.*

Proof of Proposition 1. If the information structure fails DUB, then there exist open sets $U, V \subset \Omega$ such that $U < V$ and $\mathrm{cl}(U) \cap \mathrm{cl}(V) = \emptyset$, and U is not distinguishable from V (or V is not distinguishable from U, which is analogous and hence omitted). Consider the following SCD preference with two actions: action a_1 gives a utility of 0 on all states; action a_2 gives 1 on U, -1 on V, and any number on other states such that $u(a_2, \cdot)$ is single-crossing and continuous.[34] Ω_{a_2,a_1} is not distinguishable from $\Omega^2_{a_1,a_2}$ because U is not distinguishable from V, $U \subset \Omega_{a_2,a_1}$, and $V \subset \Omega^2_{a_1,a_2}$. Proposition 5′ implies that learning fails for choice set $\{a_1, a_2\}$. *Q.E.D.*

Proof of Proposition 2. Suppose that u violates SCD. So there are actions a', a'' and states $\omega_1 < \omega_2 < \omega_3$ such that $\min\{D_{a',a''}(\omega_1), D_{a',a''}(\omega_3)\} > 0 > D_{a',a''}(\omega_2)$, where $D_{a',a''}(\omega) := u(a',\omega) - u(a'',\omega)$. We focus on priors that are only supported on $\{\omega_1, \omega_2, \omega_3\}$.

[34] Existence of such a continuous function is not trivial when the topology is not the standard Euclidean subspace topology. But one such function can be defined as follows. Firstly, define a function f' on $\mathrm{cl}(U) \cup \mathrm{cl}(V)$ as $f'(\mathrm{cl}(U)) = 1$ and $f'(\mathrm{cl}(V)) = -1$. The function is well defined because $\mathrm{cl}(U) \cap \mathrm{cl}(V) = \emptyset$. Note f' is continuous. Note also that $\mathrm{cl}(U) < \mathrm{cl}(V)$; for otherwise, $\exists u \in \mathrm{cl}(U), v \in \mathrm{cl}(V)$ s.t. $v < u$, and (v, ∞) is a neighborhood of u, which contains some point in U, contradicting $U < V$. Then, by Tietze extension theorem we can extend f' from $\mathrm{cl}(U) \cup \mathrm{cl}(V)$ to Ω as f'' such that f'' is continuous and $|f''(\Omega)| \leq 1$. Finally, define a function f over Ω such that $f(\omega) = 1$ if $\omega \leq u$ for some $u \in \mathrm{cl}(U)$; $f(\omega) = -1$ if $\omega \geq v$ for some $v \in \mathrm{cl}(V)$; and $f(\omega) = \inf_{U < \omega' \leq \omega} f''(\omega')$ otherwise. Note f is continuous and decreasing (since $|f''(\Omega)| \leq 1$), with the latter implying it is single-crossing.

Take any signal structure with strictly positive density f that satisfies MLRP (fn. 11) and DUB (e.g., normal information).[35] For any signal s and any belief $\mu \in \Delta\{\omega_1, \omega_2, \omega_3\}$,

$$
\begin{aligned}
&\mathbb{E}_{\mu_s}[D_{a',a''}(\omega)] = \sum_{i=1,2,3} \mu(\omega_i|s) D_{a',a''}(\omega_i) > 0 \\
\iff &\mu(\omega_1) f(s|\omega_1) D_{a',a''}(\omega_1) + \mu(\omega_3) f(s|\omega_3) D_{a',a''}(\omega_3) > -\mu(\omega_2) f(s|\omega_2) D_{a',a''}(\omega_2) \\
\iff & -\mu(\omega_1) \frac{f(s|\omega_1) D_{a',a''}(\omega_1)}{f(s|\omega_2) D_{a',a''}(\omega_2)} - \mu(\omega_3) \frac{f(s|\omega_3) D_{a',a''}(\omega_3)}{f(s|\omega_2) D_{a',a''}(\omega_2)} > \mu(\omega_2). \qquad (5)
\end{aligned}
$$

Under MLRP, the strictly positive functions

$$
h(s) := -\frac{f(s|\omega_1) D_{a',a''}(\omega_1)}{f(s|\omega_2) D_{a',a''}(\omega_2)} \quad \text{and} \quad g(s) := -\frac{f(s|\omega_3) D_{a',a''}(\omega_3)}{f(s|\omega_2) D_{a',a''}(\omega_2)}
$$

are decreasing and increasing, respectively, in s. Thus, given any s' and $\delta := \min\{h(s'), g(s')\} > 0$, it holds that $h(s) \geq \delta$ for all $s \leq s'$, while $g(s) \geq \delta$ for all $s \geq s'$. Consequently, for any $\varepsilon > 0$ small enough and $\mu \in \Delta\{\omega_1, \omega_2, \omega_3\}$ such that $\mu(\omega_2) = \varepsilon$ and $\mu(\omega_1) = \mu(\omega_3) = \frac{1-\varepsilon}{2}$, it holds for all signals s that

$$
\mu(\omega_1) h(s) + \mu(\omega_3) g(s) = \frac{1-\varepsilon}{2} (h(s) + g(s)) \geq \frac{1-\varepsilon}{2} \delta > \varepsilon = \mu(\omega_2). \qquad (6)
$$

Combining (6) and the equivalences leading up to (5), we have $\mathbb{E}_{\mu_s}[D_{a',a''}(\omega)] > 0$ for all s. Therefore, $\Omega_{a'',a'}$ is not distinguishable from $\Omega_{a',a''}$, and thus learning fails for choice set $\{a', a''\}$ by Proposition 5′. *Q.E.D.*

[35] Normal information satisfies DUB even when Ω is not countable. This is because normal information satisfies MLRP and pairwise unbounded beliefs. In fact, under MLRP, any sets of states $U < V$ are distinguishable from each other if and only if each $u \in U$ is distinguishable from each $v \in V$, which implies that DUB is equivalent to pairwise unbounded beliefs (since the topology on Ω is finer than the order topology). Here is the argument. The "only if" direction clearly holds even without MLRP. For the "if" direction, we establish that U is distinguishable from V. (The other case is symmetric.) By MLRP and pairwise unbounded beliefs, for any $u < v$, $f(s|u)/f(s|v)$ converges monotonically to 0 as s increases to $\sup S$. Invoking the monotone convergence theorem twice, we have that for any $\mu \in \Delta(U \cup V)$ with $\mu(V) > 0$:

$$
\begin{aligned}
\forall u \in U, \ \lim_{s \uparrow \sup S} \int_V \frac{f(s|\omega)}{f(s|u)} d\mu &= \int_V \lim_{s \uparrow \sup S} \frac{f(s|\omega)}{f(s|u)} d\mu = \infty; \\
\text{and hence } \lim_{s \uparrow \sup S} \frac{\mu(U|s)}{\mu(V|s)} &= \int_U \lim_{s \uparrow \sup S} \frac{f(s|\omega)}{\int_V f(s|\omega) d\mu} d\mu = 0.
\end{aligned}
$$

C. Proofs for Other Results

In this appendix we prove the remaining results (Lemma 1, Proposition 3, Proposition 4) within the setting considered in the main text: Ω is countable (endowed with the discrete metric); $F(\cdot|\omega)$ are absolutely continuous with respect to each to other, and so there are densities $f(\cdot|\omega) > 0$.[36]

Proof of Lemma 1. We first prove sufficiency. Take any state ω and a belief $\mu \in \Delta(\{\omega' : \omega' \leq \omega\})$ such that $\mu(\omega) > 0$. By assumption, for any $\varepsilon > 0$ there exists a positive probability set of signals $\overline{S}$ such that $\frac{f(s|\omega')}{f(s|\omega)} < \varepsilon, \forall \omega' < \omega, \forall s \in \overline{S}$. It follows that for all $s \in \overline{S}$,

$$\mu(\omega|s) = \frac{f(s|\omega)\mu(\omega)}{\sum_{\omega' \leq \omega} f(s|\omega')\mu(\omega')} = \frac{\mu(\omega)}{\mu(\omega) + \sum_{\omega' < \omega} \frac{f(s|\omega')}{f(s|\omega)}\mu(\omega')} > \frac{\mu(\omega)}{\mu(\omega) + \varepsilon}.$$

Since for any $\varepsilon > 0$ we can find a positive probability set of signals $\overline{S}$ satisfying the above inequality, we conclude that $1 \in \operatorname{Supp} \mu(\omega|\cdot)$. Since ω and μ are arbitrary, ω is distinguishable from $\{\omega' : \omega' < \omega\}$. The proof that ω is distinguishable from $\{\omega' : \omega' > \omega\}$ is analogous.

We next prove necessity assuming Ω is finite. Consider any ω and let μ be uniformly distributed over $\{\omega' : \omega' \leq \omega\}$. The distinguishability of ω from $\{\omega' : \omega' < \omega\}$ implies that for every $\varepsilon > 0$ there is a positive measure of signals $\overline{S}$ such that $\forall s \in \overline{S}$ we have $\frac{\sum_{\omega' < \omega} f(s|\omega')}{f(s|\omega)} < \varepsilon$, and so $\frac{f(s|\omega')}{f(s|\omega)} < \varepsilon$ for every $\omega' < \omega$. The argument establishing condition 2 of the lemma is analogous. *Q.E.D.*

Proof of Proposition 3. We prove that ω is distinguishable from its lower set. A symmetric argument applies for ω being distinguishable its upper set. We first prove the following claim:

Claim 2. *If g is strictly subexponential, then for any $\overline{s} > 0$ and $\varepsilon \in (0, 1)$ there is $s \geq \overline{s} + 1$ such that (i) $\sup_{k \geq 1/\overline{s}} \frac{g(s+k)}{g(s)} < \varepsilon$ and (ii) $\sup_{0 < k < 1/\overline{s}} \frac{g(s+k)}{g(s)} < 2$.*

<u>Proof of claim</u>: Suppose not, to contradiction. Then for some $\overline{s} > 0$ and $\varepsilon \in (0, 1)$, every $s \geq \overline{s} + 1$ has $k_s > 0$ such that either (i) $k_s \geq 1/\overline{s}$ and $\frac{g(s+k_s)}{g(s)} \geq \varepsilon$, or (ii) $0 < k_s < 1/\overline{s}$ and $\frac{g(s+k_s)}{g(s)} \geq 2$. Hence, for all $s \geq \overline{s} + 1$, either (i) $\frac{g(s+k_s)}{g(s)} \geq \varepsilon \geq \varepsilon^{k_s \overline{s}}$ (because $k_s \overline{s} \geq 1$), or (ii)

[36] The difficulty with extending these results to an uncountable state space is that once beliefs can put zero probability on all individual states, a set U may not be distinguishable from a set V even if every state in U is distinguishable from V. The results hold for an uncountable state space if we assume, for instance, that whenever a state u is distinguishable from a set V, there is a neighborhood around u that is distinguishable from V.

$\frac{g(s+k_s)}{g(s)} \geq 2 > \varepsilon^{k_s \bar{s}}$ (because $\varepsilon < 1$). Moreover, there is an increasing sequence $(s_i)_{i=1}^{\infty}$ such that $s_1 = \bar{s} + 1$ and for all $i > 1$, $s_i = s_{i-1} + k_{s_{i-1}}$. Note that for all i, $s_i = (\bar{s} + 1) + \sum_{j=1}^{i-1} k_{s_j}$.

First, suppose that $\sum_{i=1}^{\infty} k_{s_i} = \infty$, so that $\lim_{i\to\infty} s_i = \infty$. It holds that for all s_i, $\frac{g(s_i)}{g(\bar{s})} \geq \varepsilon^{(k_{s_{i-1}}+\cdots+k_{s_1})\bar{s}} = \varepsilon^{(s_i-\bar{s}-1)\bar{s}}$, which in turn implies that

$$(s_i - \bar{s} - 1)\bar{s}\log(\varepsilon) + \log(g(\bar{s})) \leq \log(g(s_i)). \tag{7}$$

However, since g is strictly subexponential, there is $p > 1$ such that for all large enough s_i,

$$\log(g(s_i)) \leq -(s_i)^p. \tag{8}$$

The left-hand side of inequality (7) is linear in s_i while the right-hand side of inequality (8) has exponent $p > 1$, so for large enough s_i these inequalities are in contradiction.

Next, suppose instead $\lim_{i\to\infty} s_i < \infty$. Then there is N such that for all $i \geq N$, we have $k_{s_i} < 1/\bar{s}$ and thus $\frac{g(s_{i+1})}{g(s_i)} \geq 2$. It follows that $\lim_{i\to\infty} \frac{g(s_i)}{g(s_N)} \geq \lim_{i\to\infty} 2^{i-N} = \infty$, a contradiction to the boundedness of g (being a density, g is bounded because it is uniformly continuous). $\square$

We use Claim 2 iteratively to construct a signal sequence $(s_i^*)_{i=1}^{\infty}$. Choose any $s_1^* > 0$, and for $i > 1$, choose any $s_i^* \geq s_{i-1}^* + 1$ that satisfies (i) $\sup_{k\geq 1/s_{i-1}^*} \frac{g(s_i^*+k)}{g(s_i^*)} < \frac{1}{i-1}$ and (ii) $\sup_{0<k<1/s_{i-1}^*} \frac{g(s_i^*+k)}{g(s_i^*)} < 2$. This construction is well-defined by Claim 2, with $\lim_{i\to\infty} s_i^* = \infty$.

Take any $\mu \in \Delta\{\omega' \leq \omega\}$ with $\mu(\omega) > 0$, and define $\bar{s}_i := s_i^* + \omega$. Then, for all i,

$$\forall \omega' < \omega, \quad \frac{f(\bar{s}_i|\omega')}{f(\bar{s}_i|\omega)} = \frac{g(\bar{s}_i - \omega')}{g(\bar{s}_i - \omega)} = \frac{g(s_i^* + (\omega - \omega'))}{g(s_i^*)} < 2,$$

and

$$\forall \omega' \leq \omega - (1/s_{i-1}^*), \quad \frac{f(\bar{s}_i|\omega')}{f(\bar{s}_i|\omega)} = \frac{g(s_i^* + (\omega - \omega'))}{g(s_i^*)} \leq \sup_{k\geq 1/s_{i-1}^*} \frac{g(s_i^* + k)}{g(s_i^*)} < \frac{1}{i-1},$$

and thus,

$$\frac{1 - \mu(\omega|\bar{s}_i)}{\mu(\omega|\bar{s}_i)} = \frac{\sum_{\omega'<\omega} \mu(\omega')f(\bar{s}_i|\omega')}{\mu(\omega)f(\bar{s}_i|\omega)} < \frac{1}{i-1}\frac{\sum_{\omega'\leq\omega-(1/s_{i-1}^*)} \mu(\omega')}{\mu(\omega)} + 2\frac{\sum_{\omega-(1/s_{i-1}^*)<\omega'<\omega} \mu(\omega')}{\mu(\omega)}.$$

The last expression can be taken arbitrarily small because $s_{i-1}^* \to \infty$ as $i \to \infty$. Finally, by uniform continuity of g there exists a small neighborhood of signals $\bar{S}_i$ around $\bar{s}_i$ over which the above inequality continues to hold. Therefore, $1 \in \operatorname{Supp} \mu(\omega|\cdot)$. *Q.E.D.*

The next two lemmas will be used in the proof of Proposition 4. Lemma 6 strengthens the standard result that (strict) MLRP implies that posteriors after any two signals are ordered by (strict) first-order stochastic dominance, FOSD hereafter. Lemma 7 shows that if certain histories occur in some states but not in some other state, then the latter must be distinguishable from the former.[37]

We say that signals s and s' are *equivalent* if there is a constant c such that for all ω, $f(s|\omega) = cf(s'|\omega)$. That is, equivalent signals generate the same posteriors from any prior.

Lemma 6. *Assume the signal structure f satisfies MLRP. For any full-support prior μ and non-equivalent signals $s < s'$, the posterior $\mu_{s'}$ strictly dominates the posterior μ_s in the sense of FOSD: $\forall\omega < \sup\Omega$, $\mu_s(\{\omega' \leq \omega\}) > \mu_{s'}(\{\omega' \leq \omega\})$.*

Proof. Take any signals $s < s'$ and a full-support prior μ. As is well known, MLRP implies that for all ω,

$$\mu_s(\{\omega' \leq \omega\}) \geq \mu_{s'}(\{\omega' \leq \omega\}).$$

Suppose, towards contradiction, that s and s' are not equivalent and there is $\tilde{\omega} < \sup\Omega$ such that

$$\mu_s(\{\omega' \leq \tilde{\omega}\}) = \mu_{s'}(\{\omega' \leq \tilde{\omega}\}). \tag{9}$$

As s and s' are not equivalent, there is some ω such that $\mu_s(\{\omega' \leq \omega\}) \neq \mu_{s'}(\{\omega' \leq \omega\})$, which implies $\mu_s(\{\omega' \leq \omega\}) > \mu_{s'}(\{\omega' \leq \omega\})$.

We proceed assuming $\omega < \tilde{\omega}$; a symmetric argument applies if $\omega > \tilde{\omega}$. Equation 9 implies $\mu_{s'}(\{\omega < \omega' \leq \tilde{\omega}\}) > \mu_s(\{\omega < \omega' \leq \tilde{\omega}\})$. By MLRP this means $\mu_{s'}(\{\omega' > \tilde{\omega}\}) > \mu_s(\{\omega' > \tilde{\omega}\})$:

$$\frac{\mu_{s'}(\{\omega' > \tilde{\omega}\})}{\mu_s(\{\omega' > \tilde{\omega}\})} = \frac{\int_{\{\omega' > \tilde{\omega}\}} f(s'|\omega)\,\mathrm{d}\mu(\omega)}{\int_{\{\omega' > \tilde{\omega}\}} f(s|\omega)\,\mathrm{d}\mu(\omega)} \frac{\int_\Omega f(s|\omega)\,\mathrm{d}\mu(\omega)}{\int_\Omega f(s'|\omega)\,\mathrm{d}\mu(\omega)} \geq \frac{f(s'|\tilde{\omega})}{f(s|\tilde{\omega})} \frac{\int_\Omega f(s|\omega)\,\mathrm{d}\mu(\omega)}{\int_\Omega f(s'|\omega)\,\mathrm{d}\mu(\omega)} \geq$$

$$\frac{\int_{\{\omega < \omega' \leq \tilde{\omega}\}} f(s'|\omega)\,\mathrm{d}\mu(\omega)}{\int_{\{\omega < \omega' \leq \tilde{\omega}\}} f(s|\omega)\,\mathrm{d}\mu(\omega)} \frac{\int_\Omega f(s|\omega)\,\mathrm{d}\mu(\omega)}{\int_\Omega f(s'|\omega)\,\mathrm{d}\mu(\omega)} = \frac{\mu_{s'}(\{\omega < \omega' \leq \tilde{\omega}\})}{\mu_s(\{\omega < \omega' \leq \tilde{\omega}\})} > 1,$$

which is a contradiction to Equation 9. *Q.E.D.*

Lemma 7. *State ω^* is distinguishable from the set of states V if there exists a strategy profile, a history of actions h^∞, and a finite sub-history h^m such that $\mathbb{P}(h^\infty|h^m, \omega^*) = 0$ and $\mathbb{P}(h^\infty|h^m, v)$ is bounded away from 0 across $v \in V$.*

The idea is as follows. Conditional on the finite history h^m, if the infinite history h^∞ has 0 probability in ω^* but positive probability in each $v \in V$, then there must be a sequence

[37] Both lemmas hold even if the state space is uncountable, so we write integrals instead of summations.

of "diverting" signals (i.e., that lead to histories other than h^∞) that have arbitrarily higher probability in ω^* relative to all $v \in V$. Hence, ω^* is distinguishable from V.

Proof. Suppose not. Then there exists $\mu \in \Delta(V \cup \{\omega^*\})$ with $\mu(\omega^*) > 0$ and a small $\varepsilon > 0$ such that for almost every signal s, the posterior $\mu(\omega^*|s) \leq 1 - \varepsilon$. By taking the conditional distribution of μ on V, call it $\tilde{\mu}$, and $z := \frac{\varepsilon\mu(\omega^*)}{1-\mu(\omega^*)} \in (0,1)$, we obtain for almost every s,

$$\int_V f(s|\omega)\,\mathrm{d}\tilde{\mu}(\omega) \geq z f(s|\omega^*). \tag{10}$$

Let $\mathbb{P}(a'|h^n,\omega) := \sum_{B_n} Q_n(B_n) \int_S \sigma(a'|s, h^{B_n}) f(s|\omega)\,\mathrm{d}s$ be the probability that agent n plays action a' when the state is ω and the action history is h^n. Let a^n be the action taken by agent n along h^∞ and $A^{-n} := A \backslash \{a^n\}$. It holds that:

$$\begin{aligned}
\sum_{n=m}^{\infty} \log(1 - z\mathbb{P}(A^{-n}|h^n,\omega^*)) &\geq \sum_{n=m}^{\infty} \log\left(1 - \int_V \mathbb{P}(A^{-n}|h^n,\omega)\,\mathrm{d}\tilde{\mu}(\omega)\right) \qquad \text{(using (10))}\\
&= \sum_{n=m}^{\infty} \log\left(\int_V \mathbb{P}(a^n|h^n,\omega)\,\mathrm{d}\tilde{\mu}(\omega)\right)\\
&\geq \sum_{n=m}^{\infty} \int_V \log(\mathbb{P}(a^n|h^n,\omega))\,\mathrm{d}\tilde{\mu}(\omega) \qquad \text{(by Jensen's inequality)}\\
&= \int_V \sum_{n=m}^{\infty} \log(\mathbb{P}(a^n|h^n,\omega))\,\mathrm{d}\tilde{\mu}(\omega) \qquad \text{(by Tonelli's theorem)}\\
&= \int_V \log\left(\prod_{n=m}^{\infty} \mathbb{P}(a^n|h^n,\omega)\right)\mathrm{d}\tilde{\mu}(\omega)\\
&> -\infty \qquad \text{(since } \mathbb{P}(h^\infty|h^m,\omega) \text{ is bounded across } \omega \in V\text{)}.
\end{aligned} \tag{11}$$

Below we will invoke the mathematical fact that for arbitrary sequences (S_n) and (T_n) and constant $c > 0$, if $\lim_{n\to\infty} \frac{S_n}{T_n} = c > 0$ and $\sum_n S_n < \infty$, then $\sum_n T_n < \infty$.[38] Let $S_n = -\log(1 - z\mathbb{P}(A^{-n}|h^n,\omega^*))$ and $T_n = -\log(1 - \mathbb{P}(A^{-n}|h^n,\omega^*))$. Note that $\lim_{n\to\infty} \frac{S_n}{T_n} = z \in (0,1)$ because $\lim_{x\to 0} \frac{\log(1-zx)}{\log(1-x)} = z$ and (11) implies $\lim_{n\to\infty} \mathbb{P}(A^{-n}|h^n,\omega^*) = 0$. The aforementioned mathematical fact implies that

$$\sum_{n=m}^{\infty} \log(1 - \mathbb{P}(A^{-n}|h^n,\omega^*)) > -\infty.$$

[38] For any $c' < c$ there exists N such that for all $n > N$, $S_n/T_n \geq c'$, or $T_n \leq S_n/c'$. So $\sum_n T_n \leq \sum_{n\leq N} T_n + \sum_{n>N} S_n/c' < \infty$.

As $\mathbb{P}(a^n|h^n,\omega^*) = 1 - \mathbb{P}(A^{-n}|h^n,\omega^*)$, it further follows that

$$\prod_{n=m}^{\infty} \mathbb{P}(a^n|h^n,\omega^*) > 0,$$

which contradicts $\mathbb{P}(h^\infty|h^m,\omega^*) = 0$. *Q.E.D.*

Proof of Proposition 4. Take any signal structure f that satisfies MLRP but fails DUB. We begin by constructing an SCD preference u with useful properties, and then argue that given f and u, there is inadequate learning for any full-support prior.

Since f violates DUB, take ω^* that is not distinguishable from its lower set $\{\omega : \omega < \omega^*\}$. (If ω^* is not distinguishable from its upper set, the argument is analogous.) This implies ω^* is not distinguishable from $\{\omega : \omega \leq \omega'\}$ for some $\omega' < \omega^*$. For, under MLRP, ω^* being distinguishable from its lower set is equivalent to ω^* being distinguishable from each element of the lower set, which in turn is equivalent to ω^* being distinguishable from every $\{\omega : \omega \leq \omega'\}$ with $\omega' < \omega^*$ (see fn. 35).[39]

Take any pair of actions a', a^*, and consider an SCD preference u such that $u(a',\cdot) = 0$, $u(a^*,\omega) = -1$ for $\omega < \omega^*$, $u(a^*,\omega) = 1$ for $\omega \geq \omega^*$, and all other actions are strictly dominated by both a' and a^*. Hence, an agent strictly prefers a' to a^* if and only if she ascribes probability strictly greater than $1/2$ to $\{\omega : \omega < \omega^*\}$, and is indifferent if and only if this probability is $1/2$.

Fix any full-support prior μ_0 and an equilibrium σ. It is without loss of generality to assume that for any history h, $\sigma(a'|s,h)$ is decreasing (weakly) in s. To see that, note from Lemma 6 that for non-equivalent signals $s < s'$, $\sigma(a'|s,h) > \sigma(a'|s',h)$, with either $\sigma(a'|s,h) = 1$ or $\sigma(a'|s',h) = 0$. Thus, the set of signals after which an agent randomizes must all be equivalent to each other. Since the agent is indifferent when randomizing, we can assume a constant randomization probability for all those signals without changing the public belief that is generated from a' being chosen at h.

Now assume, towards contradiction, that there is adequate learning under μ_0 and σ. Therefore, the limit belief will be almost surely supported on either $\{\omega : \omega < \omega^*\}$ or $\{\omega : \omega \geq \omega^*\}$. So there is an infinite history with a herd on a', $h^\infty \equiv (\underbrace{a_1,\ldots,a_{m-1}}_{\text{each } a_i \in \{a',a^*\}}, a', a', \ldots)$ that occurs with positive probability in state ω'. Given a finite sub-history h^n of h^∞, let $\mathbb{P}(a'|h^n,\omega) \equiv \int_S \sigma(a'|s,h^n) f(s|\omega)\,\mathrm{d}s$ be the probability that agent n plays action a' when the

[39] In this argument, MLRP is only used to rule out the case in which ω^* is distinguishable from $\{\omega : \omega \leq \omega'\}$ for all $\omega' < \omega^*$, but ω^* is not distinguishable from $\{\omega : \omega < \omega^*\}$. This property can also be guaranteed by other conditions, e.g., $\{\omega : \omega < \omega^*\}$ has a largest element.

state is ω. The probability of playing a' is decreasing in ω because $\sigma(a'|s, h^n)$ is decreasing in s and f satisfies MLRP. In particular,

$$\forall \omega \leq \omega' : \mathbb{P}(a'|h^n, \omega) \geq \mathbb{P}(a'|h^n, \omega') > 0,$$

and since h^∞ has positive probability at ω', it follows that

$$\forall \omega \leq \omega' : \prod_{n=m}^{\infty} \mathbb{P}(a'|h^n, \omega) \geq \prod_{n=m}^{\infty} \mathbb{P}(a'|h^n, \omega') > 0. \tag{12}$$

This means $\mathbb{P}(h^\infty|h^m, \omega)$ is uniformly bounded away from 0 for $\{\omega : \omega \leq \omega'\}$. On the other hand, since there is adequate learning and h^m will occur with positive probability conditional on ω^* by absolute continuity, $\mathbb{P}(h^\infty|h^m, \omega^*) = 0$. So Lemma 7 implies ω^* is distinguishable from $\{\omega : \omega \leq \omega'\}$, which is a contradiction. *Q.E.D.*

D. Overarching Probability Space and Beliefs

This appendix elaborates on the formal objects corresponding to an agent's belief and the distribution of her beliefs. To define those, we must first define a suitable probability space.

Overarching probability space. Our probability space is constructed from several elements:

1. The Markov kernel F and probability space $(\Omega, \mathcal{B}(\Omega), \mu_0)$ jointly define the probability distribution of states and signals;
2. The network structure $Q \equiv (Q_n)_{n\in\mathbb{N}}$ is independent of everything else. All Q_n's are mutually independent, respectively supported on $\{1, \ldots, n-1\}$, with the sigma-algebra being all subsets;
3. Each agent n's strategy $\sigma_n(\cdot|a_{B_n}, B_n)$ is a Markov kernel from $(A^{|B_n|}, \mathcal{A}^{|B_n|})$ to $(A, \mathcal{A})$ for each realization of neighborhood B_n.

Taken together, for the first n agents, we can define a probability space that describes the joint distribution of their neighborhoods, signals, actions, and the states. Since all these elements lie in standard Borel spaces, the Kolmogorov Extension Theorem guarantees existence of an overarching probability space $(H_\infty, \mathcal{H}_\infty, \mathbb{P})$ that is consistent with each finite probability space (i.e., up to each agent n).

Beliefs. Given this overarching probability space, the interim belief (i.e., the belief after observing her neighbors and their actions, but before observing her private signal) of agent n is $\mathbb{P}(\cdot|a_{B_n}, B_n)$, and the posterior belief of agent n is $\mathbb{P}(\cdot|a_{B_n}, B_n, s_n)$. These beliefs are well defined because, as a countable product of standard Borel spaces, the overarching probability space is a standard Borel space, and hence there exist regular conditional probabilities (Durrett, 2019, Theorem 4.1.17).

Distribution of beliefs. The interim belief of agent n, μ_n, as a regular conditional probability, can be regarded as a measurable function from $(H_\infty, \mathcal{H}_\infty, \mathbb{P})$ to $(\Delta\Omega, \mathcal{B}(\Delta\Omega))$; see Crauel (2002, Remark 3.20). As Ω is a Polish space, so is $\Delta\Omega$. We define agent n's distribution of (interim) beliefs, φ_n, as the push-forward measure of μ_n. Hence, $\varphi_n \in \Delta\Delta\Omega$ since it is by definition a Borel probability measure on $\Delta\Omega$.