

Beyond Unbounded Beliefs: How Preferences and Information Interplay in Social Learning*

Navin Kartik SangMok Lee Tianhao Liu Daniel Rappoport[†]

April 9, 2024

Abstract

When does society eventually learn the truth, or take the correct action, via observational learning? In a general model of sequential learning over social networks, we identify a simple condition for learning dubbed *excludability*. Excludability is a joint property of agents' preferences and their information. We develop two classes of preferences and information that jointly satisfy excludability: (i) for a one-dimensional state, preferences with single-crossing differences and a new informational condition, directionally unbounded beliefs; and (ii) for a multi-dimensional state, intermediate preferences and sub-exponential location-shift information. These applications exemplify that with multiple states "unbounded beliefs" is not only unnecessary for learning, but incompatible with familiar informational structures like normal information. Unbounded beliefs demands that a single agent can identify the correct action. Excludability, on the other hand, only requires that a single agent must be able to displace any wrong action, even if she cannot take the correct action.

1. Introduction

This paper concerns the classic sequential observational or social learning model initiated by [Banerjee \(1992\)](#) and [Bikhchandani, Hirshleifer, and Welch \(1992\)](#). There

*We thank Nageeb Ali, Marina Halac, Ben Golub, Andreas Kleiner, Elliot Lipnowski, José Montiel Olea, Xiaosheng Mu, Harry Pei, Jacopo Peregó, Evan Sadler, Lones Smith, Peter Sørensen, the Editor, and three anonymous referees 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.

[†]Kartik and Liu are at Columbia University; Lee is at Washington University in St. Louis; and Rappoport is at the University of Chicago.

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). Agents act in sequence, each receiving her own private information about the state and observing some subset 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)?

One would anticipate that whether there is social learning depends on the combination of agents' preferences and their information structure. But, at least for finite action sets, economists have largely emphasized the latter dimension alone.¹ The reason is inextricably tied to focusing on models with two states. With only two states, there is social learning given any (nontrivial) preferences if and only if there is learning for all preferences. For, with two states, even the former requires private signals/beliefs to be *unbounded* (Smith and Sørensen, 2000; Acemoglu, Dahleh, Lobel, and Ozdaglar, 2011). 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.

With multiple—i.e., more than two—states, unbounded beliefs still characterizes learning for all preferences (Arieli and Mueller-Frank, 2021).² However, it is now a very demanding condition. Consider, for instance, the canonical example of *normal information*: the state is $\omega \in \Omega \subset \mathbb{R}$ and agents' signals are drawn independently from a normal distribution with mean ω and fixed variance. With only two states, there is unbounded beliefs because a very high signal makes one arbitrarily convinced of the high state, while a very low signal makes one arbitrarily convinced of the low state. But with multiple states, normal information fails unbounded beliefs: given any full-support prior, there is an upper bound on

¹Unless noted otherwise, our introduction should be understood as referring to the canonical sequential social learning model with a finite action set, homogeneous preferences, and no direct payoff externalities. It is well recognized that variations in those aspects can also matter for social learning; see for example, Lee (1993) on infinite action spaces, Avery and Zemsky (1998) and Eyster, Galeotti, Kartik, and Rabin (2014) on endogenous prices or congestion costs, and Goeree, Palfrey, and Rogers (2006) on heterogeneous preferences.

²Arieli and Mueller-Frank (2021, Theorem 1) refer to the condition as “totally unbounded beliefs”. They establish their result for a complete network, i.e., when each agent observes the actions of all predecessors. A by-product of our analysis is to establish it for general networks (Corollary 1 in Section 3).

how certain one can become about any non-extremal state based on observing one signal.³ Is social learning doomed with multiple states for familiar information structures like normal information?

Our paper shows that the answer is no. With multiple states, whether society eventually learns to take the correct action depends on the interplay of preferences and information. Crucially, learning can obtain under standard preferences with familiar information structures that fail unbounded beliefs. [Figure 1](#) illustrates an example of normal information with state space $\Omega = \{1, 2, 3\}$, action set $A = \{a_1, a_2\}$, and a uniform prior μ_0 . The failure of unbounded beliefs is reflected in the set of posteriors, represented by the black curve, being bounded away from state 2’s vertex. For concreteness, suppose that each agent observes all predecessors’ actions. In [Figure 1a](#), preferences violate single crossing—defined formally in [Section 4](#)—because action a_1 is optimal in both states 1 and 3, whereas a_2 is optimal in state 2. Here, learning fails: since action a_1 is optimal after any signal the first agent receives, society is stuck with all agents taking a_1 . By contrast, in [Figure 1b](#), agents have single-crossing preferences; specifically an agent who takes action a_i gets the quadratic-loss utility $-(i - \omega)^2$. Now, at any belief at which learning the state would be useful (i.e., a belief that puts positive probability on both state 1, where a_1 is optimal, and either state 2 or 3, where a_2 is optimal), no single action is optimal after all signals. This property yields social learning; see [Theorem 1](#) in [Section 3](#).

Excludability. Our paper develops a simple joint condition on information and preferences, which we call *excludability*, that is not only sufficient for social learning on general observational networks (satisfying a mild condition known as [expanding observations](#)), but in a sense also necessary; see [Theorem 2](#) in [Section 3](#).

Roughly speaking, excludability requires that for each pair of actions, a and a' , a single agent must be able to receive a signal that makes her arbitrarily convinced that a is better than a' , no matter which (full-support) belief she starts with. Put differently, information must be able to distinguish the set of states in which a is better than a' from the set in which a' is better than a . Excludability implies that society can never get stuck on a wrong action: if an action is suboptimal at the true state, then some agent will receive a private signal convincing her not to take that

³So binary states is special because all states are extreme states. There is nothing exceptional about normal information violating unbounded beliefs; see [Remark 2](#) in [Section 3](#).

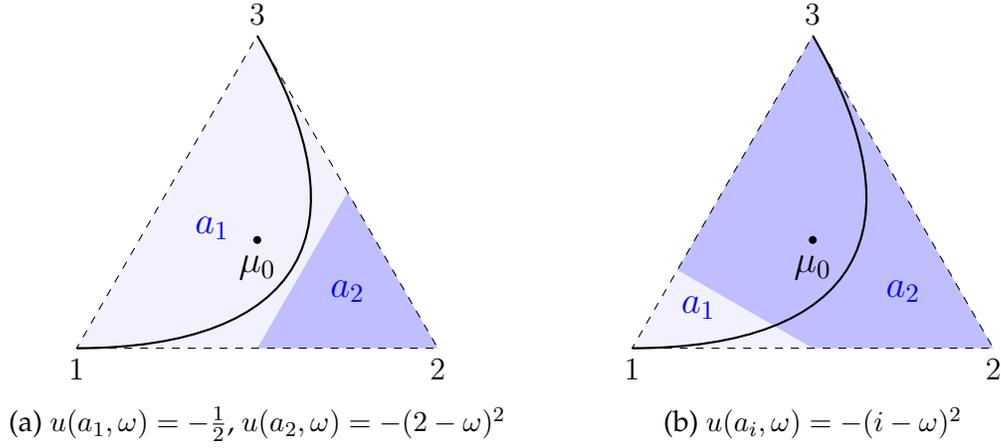


Figure 1: Belief simplex for state space $\Omega = \{1, 2, 3\}$. The curve depicts the set of posteriors for a single agent under normal information with prior μ_0 . The action set is $A = \{a_1, a_2\}$ and each agent's utility is $u(a, \omega)$. The shaded regions depict optimal actions under uncertainty.

action. We establish that this property of *displacing wrong actions* leads to social learning. Notably, an agent can displace wrong actions even if she cannot take the correct action, i.e., the optimal action at the true state. (See Figure 2 in Section 3 for a concrete example.) We view the distinction of social learning arising from the individual capacity to displace wrong actions rather than to discover the correct action as a key insight; this distinction² cannot be seen with only two states, where the two notions are equivalent.

Excludability provides a useful perspective on existing ideas in the literature. For instance, as detailed in Section 3, an information structure yields excludability for all preferences if and only if that information structure has unbounded beliefs. But more importantly, we can use excludability to deduce weaker informational conditions that yield social learning for canonical classes of preferences.⁴

Single-crossing preferences. Our leading application of excludability is to preferences with *single-crossing differences* (SCD). Here we show that learning obtains when the information structure satisfies *directionally unbounded beliefs* (DUB). SCD is a familiar property (Milgrom and Shannon, 1994) that is widely assumed in eco-

⁴ Although this approach of obtaining more tenable conditions by restricting preferences to some broad class is novel to social learning, it is classical in other areas of economics. For instance, first-order stochastic dominance is weakened to second-order by restricting to concave (and increasing) utility functions.

nomics: it captures settings in which there are no preference reversals as the state increases. By contrast, DUB appears to be a new condition on information structures, although Milgrom (1979) utilizes a related property in the context of auction theory. Like SCD, DUB is formulated for a (totally) 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 . Crucially, no signal need make one arbitrarily certain about ω (unlike unbounded beliefs). For the normal information structure discussed earlier, requirements (i) and (ii) are met for any state by arbitrarily high and arbitrarily low signals, respectively.

Proposition 1 in Section 4 shows that SCD preferences and DUB information are jointly sufficient for excludability, and hence learning. For a direct intuition on the SCD-DUB interplay, consider normal information again. There are preferences (like those in Figure 1a) under which society can get stuck at some belief at which agents are taking an incorrect action, but only a strong signal about an intermediate state would change the action—alas no such signal is available. However, under SCD preferences (like those in Figure 1b), if knowing that the state is some intermediate ω would change the action, then so would knowing that the state is at least ω or at most ω . Normal information, or more generally DUB, guarantees that there are strong signals approximating such knowledge.

Intermediate preferences. Our second application in Section 4 is to *intermediate preferences* in multidimensional spaces (Grandmont, 1978), where the state is $\omega \in \mathbb{R}^d$ and the action is $a \in \mathbb{R}^d$. These subsume both constant-elasticity-of-substitution preferences common in many areas of economics and Euclidean preferences invoked in political economy and communication/delegation models.

Using excludability, we show that social learning obtains under intermediate preferences so long as information is given by a *subexponential location-shift family*. Location-shift families are widely-used information structures: for some density $g : \mathbb{R}^d \rightarrow \mathbb{R}^d$, the signal distribution in any state ω is given by $g(s - \omega)$. Loosely, the subexponential condition requires that the tail of g must be thin enough, eventually decreasing faster than an exponential rate. We establish that this thin-tails property combined with intermediate preferences yields excludability. Notably, multidimensional normal information (i.e., normally distributed signals with mean equal to the state and some fixed covariance matrix) satisfies the subexponential

requirement.

Methodology. A significant contribution of our paper is also methodological. We develop an approach to tackle learning, and more generally, asymptotic social welfare with multiple states in general observational networks. [Theorem 1](#) in [Section 3](#) is the backbone by which we tie learning to excludability. [Theorem 1](#) reduces the complex dynamic problem of social learning in networks to a much simpler “static” problem. The theorem says that there is learning if and only if every *stationary belief* has *adequate knowledge*. A stationary belief is one at which there is an action that is optimal no matter an agent’s signal, and an adequate-knowledge belief is one at which there is an action that is optimal no matter the state in the belief’s support. Excludability is a simple sufficient—and necessary, in a sense explained later—condition for all stationary beliefs to have adequate knowledge.

[Theorem 1](#) itself is a consequence of [Theorem 3](#) in [Section 5](#), which provides a welfare lower bound even when learning fails. The theorem roughly says that for any preferences 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. [Theorem 3](#) implies that learning obtains when the cascade utility equals the utility obtained from taking the correct action in each state, which leads to [Theorem 1](#).

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 with binary states, [Smith and Sørensen \(2000\)](#) show that, given any nontrivial preferences, there is learning if and only if beliefs are unbounded. For the complete network but with multiple states, [Arieli and Mueller-Frank \(2021\)](#) 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 preferences.⁵ The approach of both [Smith and Sørensen \(2000\)](#) and [Arieli and Mueller-Frank \(2021\)](#) rests on the social belief—an agent’s belief based on observing her predecessors’ actions, before observing her own

⁵The 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.

signal—being a martingale in the complete network.

Gale and Kariv (2003) and Çelen and Kariv (2004) depart from the complete network, noting that martingale methods now fail. Both these papers also depart from the canonical setting in other ways, however: in Gale and Kariv (2003) agents choose actions repeatedly, while in Çelen and Kariv (2004) private signals are not independent conditional on the true state. Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) provide a general treatment of observational networks in an otherwise classical setting. But they only allow for binary states and binary actions. They introduce the condition of expanding observations, explaining that this property of the network is necessary for learning. They establish that it is also sufficient for learning with unbounded beliefs. Building on Banerjee and Fudenberg (2004), a key contribution of Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) is to use a welfare *improvement principle* to deduce learning; 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.

The analysis in both Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) and Lobel and Sadler (2015) relies on their binary-state binary-action structure.⁶ We believe 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.

Substantively, our focus on multiple states and actions allows us to shed light on how preferences and information jointly shape social learning. As already noted, their interplay in determining learning has not received attention in the prior literature because of its focus on binary states. The only exception we are aware of is Arieli and Mueller-Frank (2021, Theorem 3), discussed in Section 3; their result assumes a special utility function and is only for the complete network.

⁶Banerjee and Fudenberg (2004) and Smith and Sørensen (2020) consider “unordered” random sampling models that also only allow for binary states and actions.

2. Model

There is a countable state space Ω , endowed with the discrete topology, and standard Borel spaces of actions A and signals S . We allow each of these three sets to be finite or infinite. 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 to some reference measure that is mutually absolutely continuous with every $F(\cdot|\omega')$. Without further loss of generality we assume $f(\cdot|\cdot) > 0$, so that no signal rules out any state.

The game. At the outset, a state ω is drawn from a common prior probability mass function $\mu_0 \in \Delta\Omega$.⁷ Then, an infinite sequence of agents, indexed by $n = 1, 2, \dots$, sequentially select actions. An agent n observes both a private signal s_n drawn from $f(\cdot|\omega)$ and the actions of some subset of her predecessors $B_n \subseteq \{1, 2, \dots, n-1\}$, and then chooses her action $a_n \in A$. Agents' private signals are drawn independently conditional on the state, and 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 $\{1, 2, \dots, n-1\}$, 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 realized 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 $(a_k)_{k \in B_n}$.⁸ Let \mathcal{I}_n denote the set of all possible information sets for agent n . A strategy for agent n is a (measurable) function $\sigma_n : \mathcal{I}_n \rightarrow \Delta A$.

All agents are expected utility maximizers and have common preferences that depend only on their own action and the state, represented by the utility function $u : A \times \Omega \rightarrow \mathbb{R}$. We assume that utility is bounded: there is $\bar{u} \geq 0$ such that

⁷For any topological space X , ΔX denotes the set of Borel probability measures over X .

⁸While we assume that each agent observes the identities of her neighbors as well as their chosen actions, the **Conclusion** explains how our analysis extends to various cases of "random sampling" in which neighbors' identities are not observed. Our analysis also applies if agents receive arbitrary information about their predecessors' realized neighborhoods.

$$|u(\cdot, \cdot)| \leq \bar{u}.$$

We study the Bayes Nash equilibria—hereafter simply equilibria—of this game. We assume that for every belief there is an optimal action, so that an equilibrium exists.⁹

Remark 1. **Appendix A** describes a more general setting in which our main results are proved. For example, Ω can be a closed subset of \mathbb{R} and each $u(a, \cdot)$ piecewise continuous with A finite. We also do not require the signal distributions to be mutually absolutely continuous.

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 σ , $\mathbb{E}_{\sigma, \mu_0}[u_n] \rightarrow u^*(\mu_0)$. 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.¹⁰ We say there is *inadequate learning* if adequate learning fails.¹¹

We will also be interested in situations in which agents choose from a subset of actions, referred to as a *choice set*.¹² We say that there is *(in)adequate learning for a choice set* $\tilde{A} \subseteq A$ if there is (in)adequate learning when agents are restricted to choose from actions in \tilde{A} .

⁹Existence of optimal actions is assured under standard assumptions, e.g., if 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 . Hence, we could just as well adopt (weak) Perfect Bayesian equilibrium or refinements.

¹⁰Our notion of adequate learning is different from [Arieli and Mueller-Frank's \(2021\)](#), who require learning for all utility functions. Following [Aghion, Bolton, Harris, and Jullien \(1991\)](#), we use "adequate" to signify that learning the state precisely is not necessary when some action is optimal in multiple states.

¹¹That we deem learning to be inadequate if there is some equilibrium in which learning fails, rather than in every equilibrium, is innocuous given that there is no strategic interaction (cf. [fn. 9](#)). On the other hand, the issue of whether learning fails at every prior rather than only at some priors is substantive. We return to this issue in our [Conclusion](#).

¹²We restrict attention to choice sets such that for every belief there is an optimal action.

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

$$\forall K \in \mathbb{N} : \lim_{n \rightarrow \infty} Q_n(B_n \subseteq \{1, \dots, K\}) = 0. \quad (1)$$

The reason is that 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 at most actions a_1, \dots, a_K . In that event, the agent cannot do better than choosing her action based on only $K + 1$ signals.

Accordingly, we assume expanding observations. Leading examples of network structures with expanding observations include: (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 $k \in \{1, \dots, n - 1\}$).

3. Characterizations of Learning

3.1. Stationary Beliefs and Adequate Knowledge

The key to all our results on learning is [Theorem 1](#) below, which simplifies the question of adequate learning to a “one-shot updating” property of beliefs. To state that result, we require two concepts concerning the value of information.

For any belief $\mu \in \Delta\Omega$, let $c(\mu) := \arg \max_{a \in A} \mathbb{E}_\mu[u(a, \omega)]$ denote the set of optimal actions under that belief. 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 μ_s , we say that belief μ is *stationary* if there is $a \in c(\mu)$ such that $a \in c(\mu_s)$ for μ -a.e. signal s . We say that belief μ has *adequate knowledge* if there is $a \in c(\mu)$ such that $a \in c(\omega)$ for all $\omega \in \text{Supp } \mu$. So a belief is stationary if an agent holding that belief does not benefit from observing a signal from the given information structure.¹³ On the other hand, a belief has adequate knowledge if the agent would not benefit from observing a signal from *any* information structure, in particular learning the state.

¹³Some readers may find it helpful to note that in their setting, [Smith and Sørensen \(2000\)](#) refer to stationary beliefs as “cascade beliefs”.

Any adequate-knowledge belief, such as a belief that puts probability one on a single state, is stationary. In general, there can be stationary beliefs without adequate knowledge, as seen in [Figure 1a](#).

Theorem 1. *There is adequate learning if and only if all stationary beliefs have adequate knowledge.*

[Theorem 1](#) provides a characterization of adequate learning that holds regardless of the observational network structure, given our maintained assumption of expanding observations. Its “only if” direction is straightforward because our notion of learning considers all priors: if the prior is stationary and has inadequate knowledge, then society is stuck with all agents taking the prior-optimal action even though it is suboptimal in some states. More important and subtle is the theorem’s “if” direction. It is inspired by earlier results, particularly [Arieli and Mueller-Frank \(2021, Lemma 1\)](#) and [Lobel and Sadler \(2015, Theorem 1\)](#), but the logic in the current general setting of arbitrary networks and multiple states and actions is novel. We defer this logic to [Section 5](#), instead turning now to how we can build on [Theorem 1](#) for a more practicable characterization of learning. In particular, we seek a more transparent condition on the combinations of preferences and information that yield adequate learning.

3.2. Excludability

A key notion is whether information allows an agent to become arbitrarily sure about a subset of states Ω' relative to another subset Ω'' . To make that precise, let $\mu_s(\Omega')$ denote the posterior on states Ω' induced by belief μ and signal s , and $\Pr_\mu(S')$ be the probability of signal set S' induced by belief μ .

Definition 1. A set Ω' is *distinguishable* from another set Ω'' if for any $\varepsilon > 0$ and $\mu \in \Delta(\Omega' \cup \Omega'')$ with $\mu(\Omega') > 0$, it holds that $\Pr_\mu(s : \mu_s(\Omega') > 1 - \varepsilon) > 0$.

Note that Ω' is distinguishable from Ω'' if and only if every $\omega \in \Omega'$ is distinguishable from Ω'' . Moreover, if Ω' is distinguishable from Ω'' , then every subset of Ω' is distinguishable from every subset of Ω'' . The following observation essentially reinterprets distinguishability directly in terms of the signal structure rather than posteriors.

Lemma 1. Ω' is distinguishable from Ω'' if for every $\omega' \in \Omega'$ and $\varepsilon > 0$, there is a positive-probability set of signals S' such that

$$\forall \omega'' \in \Omega'', \forall s \in S' : \frac{f(s|\omega'')}{f(s|\omega')} < \varepsilon.$$

Conversely, this condition is also necessary if Ω'' is finite.

We emphasize that the set S' in the lemma cannot depend on $\omega'' \in \Omega''$; for Ω' to be distinguished from Ω'' , each $\omega' \in \Omega'$ must be distinguished from all $\omega'' \in \Omega''$ simultaneously. Consider the example of *normal information*: $\Omega \subset \mathbb{R}$ and signals are normally distributed on \mathbb{R} with mean ω and fixed variance. When $\Omega = \{1, 2, 3\}$, state 2 is distinguishable from 1 because $f(s|1)/f(s|2) \rightarrow 0$ as $s \rightarrow \infty$, and state 2 is distinguishable from 3 because $f(s|3)/f(s|2) \rightarrow 0$ as $s \rightarrow -\infty$. But state 2 cannot be distinguished from both 1 and 3 simultaneously, because $\min\{f(s|1)/f(s|2), f(s|3)/f(s|2)\}$ is bounded away from 0.

Distinguishability of each state from its complement 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 with multiple states, unbounded beliefs is incompatible with familiar information structures.

Remark 2. Under any monotone likelihood ratio property (MLRP) information structure, no state ω is distinguishable from $\{\omega', \omega''\}$ with $\omega' < \omega < \omega''$.¹⁴ Consequently, if $|\Omega| > 2$, unbounded beliefs fails under the MLRP.

Fortunately, learning only requires certain subsets of states to be distinguished from each other. For any two actions a_1 and a_2 , let the preferred set $\Omega_{a_1, a_2} := \{\omega : u(a_1, \omega) > u(a_2, \omega)\}$ be the set of states in which a_1 is strictly preferred to a_2 .

Definition 2. A utility function and an information structure jointly satisfy *excludability* if for every a_1 and a_2 , Ω_{a_1, a_2} is distinguishable from Ω_{a_2, a_1} .

Excludability is a joint condition on preferences and information. It requires that for any pair of actions, a single agent can become arbitrarily certain that one action is strictly better than the other, starting from any belief that does not exclude

¹⁴For ordered state and signals spaces, the MLRP holds if $\forall s' > s$ and $\forall \omega' > \omega$, $f(s|\omega')/f(s|\omega) \leq f(s'|\omega')/f(s'|\omega)$.

that event. Since excludability is defined using preferred sets, it is straightforward to deduce which sets must be distinguishable for any given preferences; [Lemma 1](#) then provides a set of likelihood-ratio conditions on the information structure, without reference to beliefs.

Unbounded beliefs implies excludability for any preferences. Conversely, if unbounded beliefs fails, then there is some state ω^* that is not distinguishable from its complement, and excludability fails when preferences are such that for some a_1 and a_2 , $\Omega_{a_1, a_2} = \{\omega^*\}$ while $\Omega_{a_2, a_1} = \Omega \setminus \{\omega^*\}$. Hence, excludability for all preferences is equivalent to unbounded beliefs. But with multiple states, excludability can be substantially weaker for any given (class of) preferences, as developed in [Section 4](#).¹⁵ This matters because:

Theorem 2. *Excludability implies adequate learning for every choice set. If excludability fails and the number of states is finite, then there is inadequate learning for some choice set.*

(See [Theorem 2'](#) in the appendix for a more general version of [Theorem 2](#) that does not require finiteness in the second statement. Hereafter, for brevity, we leave it as implicit that it is [Theorem 2'](#) rather than [Theorem 2](#) we are invoking when discussing the necessity of excludability for learning in an infinite state space.)

Excludability is sufficient for adequate learning because it ensures that wrong actions can always be “displaced”, which by [Theorem 1](#) is the key to social learning. More precisely, excludability guarantees that, no matter the choice set, all stationary beliefs have adequate knowledge. Suppose a belief μ has inadequate knowledge, so that $c(\mu) \neq c(\omega^*)$ for some state $\omega^* \in \text{Supp } \mu$. (For simplicity, assume $c(\mu)$ and $c(\omega^*)$ are singletons.) Excludability implies that preferred set $\Omega_{c(\omega^*), c(\mu)}$ is distinguishable from $\Omega_{c(\mu), c(\omega^*)}$. Hence, with positive probability, an agent who starts with belief μ will obtain a posterior that puts arbitrarily large probability on $\Omega_{c(\omega^*), c(\mu)}$ relative to $\Omega_{c(\mu), c(\omega^*)}$, in which event she strictly prefers $c(\omega^*)$ to $c(\mu)$. Consequently, μ is not stationary.

We highlight that excludability does not guarantee that a wrong action can always be displaced by the correct action. In other words, even though excludability guarantees that given any wrong action—say, $c(\mu)$ when the true state is ω^* —a

¹⁵With only two states, $\Omega = \{\omega_1, \omega_2\}$, excludability under any given nontrivial preferences is equivalent to unbounded beliefs. (Nontrivial means that no action is optimal at all states.) For, there must be actions a_1 and a_2 such that $\Omega_{a_1, a_2} = \{\omega_1\}$ and $\Omega_{a_2, a_1} = \{\omega_2\}$; excludability requires these sets to be mutually distinguishable, which is unbounded beliefs.

single agent can receive a signal convincing her that $c(\mu)$ is worse than the correct action $c(\omega^*)$, there may be no signal that leads the agent to take $c(\omega^*)$. When there are two states and finite actions, always being able to displace a wrong action and always being able to take the correct action are equivalent, as they both reduce to unbounded beliefs. But more generally, it is displacing wrong actions that is fundamental for learning.

To illustrate the point concretely, consider the example depicted in [Figure 2](#). There are three states and three actions, $\Omega = A = \{1, 2, 3\}$. The signal structure and preferences are detailed in the figure’s caption. The correct action in each state ω is $a = \omega$. Importantly, unbounded beliefs fails yet there is excludability.¹⁶ Let agent n ’s *social belief* be her belief about the state given only the history of her neighbors’ actions, prior to observing her own private signal. When each agent observes all predecessors’ actions, [Figure 2](#) shows two representative numerically-simulated paths of social beliefs given the true state $\omega = 2$. The social belief starts at the prior, marked by a star in the figure, and then evolves as agents take actions, as indicated by either of the arrowed paths. There is a range of beliefs, shaded in grey, such that for any social belief in that range no signal can lead an agent to take the correct action 2. As the prior is in this range, *the first agent necessarily takes a wrong action*: either 1 (which occurs in the red path) or 3 (the blue path). Nevertheless, even though no agent can take the correct action 2 for a while, society never gets stuck at a wrong action: given that an agent’s predecessor chose $a \in \{1, 3\}$, there are signals (very high if $a = 1$ and very low if $a = 3$) that convince the agent that a is worse than the correct action 2, and hence the agent will not take action a . At some point, after enough switching between actions 1 and 3, the social belief is driven outside the grey region and it becomes possible for an agent to take the correct action 2. Eventually, society settles on that action.

Turning to necessity in [Theorem 2](#): for a fixed choice set, all stationary beliefs can have adequate knowledge (and hence there is adequate learning, by [Theorem 1](#)) even absent excludability. But when excludability fails, there is some preferred set Ω_{a_1, a_2} that cannot be distinguished from Ω_{a_2, a_1} . If Ω is finite, this means

¹⁶Unbounded beliefs fails because under normal information the middle state is not distinguishable from its complement. Excludability can be verified by checking distinguishability of the preferred sets for each pair of actions; alternatively, we note that the preferences satisfy [single-crossing differences](#) (SCD), and as explained in [Subsection 4.1](#), SCD and normal information imply excludability.

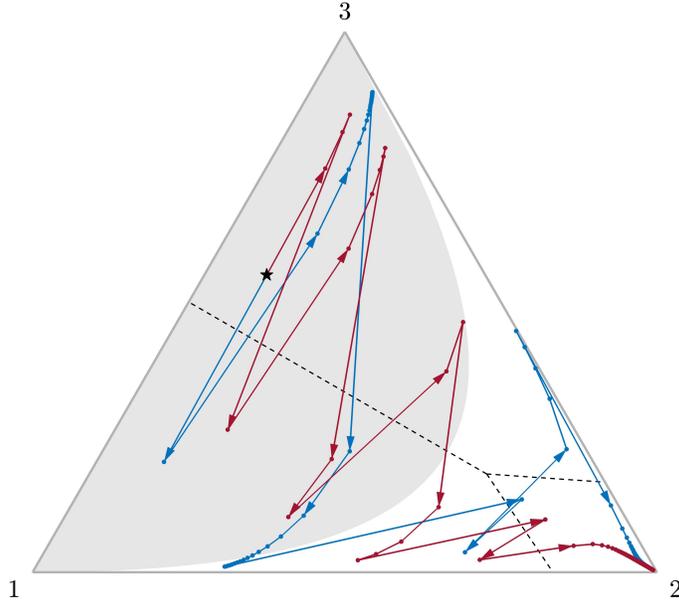


Figure 2: Two simulated social belief paths—one in red and one in blue—in a complete network. There are three states labeled 1, 2, 3, and there is normal information (with standard deviation 1.2). There are three actions with respective state-contingent utilities $(1, 0, -0.3)$, $(0, 0.2, 0)$, and $(-0.3, 0, 1)$. The optimal action under uncertainty is delineated by the dashed lines. The true state is 2, and society starts with the prior $(0.35, 0.1, 0.55)$, marked by the black star. The grey shaded region indicates beliefs at which no single signal can lead to state 2’s correct action. On each path, a dot represents the social belief after an agent has acted, and arrows indicate the sequencing.

that when the choice set is $\{a_1, a_2\}$, a belief that puts small probability on Ω_{a_1, a_2} relative to Ω_{a_2, a_1} is stationary and has inadequate knowledge. Hence, **Theorem 1** implies that excludability is necessary for learning when we seek learning for all choice sets. The following example illustrates these points using an infinite action set for convenience.

Example 1. Consider $\Omega = \{0, 1\}$, $A = [0, 1]$, and $u(a, \omega) = -(a - \omega)^2$. This is an example of “responsive preferences” (Lee, 1993; Ali, 2018). Fix any nontrivial signal structure and any observational network structure satisfying expanding observations.

Adequate learning obtains by **Theorem 1**, because the only stationary beliefs have certainty on one of the two states. For, given any nondegenerate belief, with positive probability the posterior-optimal action will be different from the prior-optimal action, as the uniquely optimal action equals the posterior expected state.

However, excludability is equivalent to the signal structure having unbounded beliefs, as for any $a_1 < a_2$, $\Omega_{a_1, a_2} = \{0\}$ and $\Omega_{a_2, a_1} = \{1\}$. So excludability is not necessary for adequate learning at choice set A . But absent excludability there is *inadequate* learning at any non-singleton finite choice set. For, there is then some state such that any prior that puts probability close to 1 on that state will be stationary, but this prior has inadequate knowledge. \square

The choice-set variation required by [Theorem 2](#) comes “for free” when we seek an informational condition that ensures learning for a broad-enough class of preferences. Specifically, it is sufficient that for any utility function in the class and any choice set, there is another utility function that is identical on that set but makes all other actions dominated. Since the class of all preferences has this property, and excludability for all preferences is equivalent to the information structure having unbounded beliefs, [Theorem 2](#) immediately implies:

Corollary 1. *An information structure yields adequate learning for all preferences if and only if it has unbounded beliefs.*

This corollary extends results from the prior literature, which are either for the complete network ([Arieli and Mueller-Frank, 2021](#), Theorem 1) or general networks but with only two states ([Acemoglu, Dahleh, Lobel, and Ozdaglar, 2011](#), Theorem 2).

To our knowledge, the only prior exception to unbounded beliefs driving learning with a discrete action space is the interesting example of [Arieli and Mueller-Frank \(2021, Theorem 3\)](#). They consider the complete network and a special utility function, which they call “simple utility”, in which the payoff is 1 if the action matches the state and 0 otherwise. For this case, they show that *pairwise distinguishability*—for any pair of states, each is distinguishable from the other—is sufficient for learning. This result also follows from [Theorem 2](#); indeed, the theorem implies that learning obtains for general observational networks. For, under simple utility, the preferred sets for actions $a_1 \neq a_2$ are just $\{a_1\}$ and $\{a_2\}$, which means excludability is equivalent to pairwise distinguishability.

4. Applications

Excludability permits a study of informational conditions that assure adequate learning for broad and widely-used classes of preferences. This section presents

two such applications: one with a one-dimensional state, and one with a multi-dimensional state.

4.1. Learning in a One-Dimensional World

In this subsection we assume a totally ordered state space: for simplicity, $\Omega \subset \mathbb{R}$. A function $h : \Omega \rightarrow \mathbb{R}$ is *single crossing* if either: (i) for all $\omega < \omega'$, $h(\omega) > 0 \implies h(\omega') \geq 0$; or (ii) for all $\omega < \omega'$, $h(\omega) < 0 \implies h(\omega') \leq 0$. That is, a single-crossing function switches sign between strictly positive and strictly negative at most once.

Definition 3. Preferences represented by $u : A \times \Omega \rightarrow \mathbb{R}$ 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 notions in [Milgrom and Shannon \(1994\)](#) and [Athey \(2001\)](#), but, following [Kartik, Lee, and Rappoport \(2023\)](#), the formulation is without an order on A .¹⁷ Ignoring indifferences, SCD requires that the preference over any pair of actions can only flip once as the state changes monotonically. SCD is widely satisfied in economic models; in particular, it is assured by supermodularity of u .

The key informational condition is that of distinguishing upper and lower sets from each other. More precisely, we require that for any ω , $\{\omega' : \omega' \geq \omega\}$ and $\{\omega' : \omega' < \omega\}$ are distinguishable from each other, and $\{\omega' : \omega' > \omega\}$ and $\{\omega' : \omega' \leq \omega\}$ are distinguishable from each other. But since a set Ω' is distinguishable from Ω'' if and only if each $\omega \in \Omega'$ is distinguishable from Ω'' , we can simplify as follows.

Definition 4. An information structure has *directionally unbounded beliefs* (DUB) if every ω is distinguishable from $\{\omega' : \omega' < \omega\}$ and also from $\{\omega' : \omega' > \omega\}$.

Crucially, DUB does not require any state ω to be distinguishable from any subset of states containing both a higher and a lower state than ω . Rather, using [Lemma 1](#), we can view DUB as only requiring that for any state ω , there are signals that are arbitrarily more likely in ω relative to all $\omega' < \omega$, and also other signals that are arbitrarily more likely in ω relative to all $\omega' > \omega$.

A leading example of DUB information is normal information. More generally, for any MLRP information structure, DUB can be easily checked because it

¹⁷SCD is equivalent to there existing some order on A with respect to which [Athey's \(2001\)](#) “weak single-crossing property of incremental returns” holds.

reduces to pairwise distinguishability.¹⁸ We note that when A is finite, pairwise distinguishability is inescapable (even without the MLRP) for adequate learning in any rich-enough class of preferences.¹⁹

Our main result in this subsection is:

Proposition 1. *If preferences have SCD and the information structure has DUB, then there is adequate learning. Conversely, if the information structure violates DUB and there are at least two actions, then there are SCD preferences for which there is inadequate learning.*

The result says that not only is DUB a sufficient informational condition for adequate learning under any SCD preferences, but it also necessary to assure learning for all SCD preferences.

Here is the logic for sufficiency. Recall that $\Omega_{a,a'}$ denotes the states in which action a is strictly preferred to a' . SCD implies non-reversal of strict preferences: for any a and a' , either $\inf \Omega_{a,a'} \geq \sup \Omega_{a',a}$ or $\inf \Omega_{a',a} \geq \sup \Omega_{a,a'}$. DUB says that every upper (resp., lower) set of states and its strict lower (resp., strict upper) set are distinguishable from each other. Therefore, SCD and DUB together guarantee excludability, and so [Proposition 1](#)'s first statement follows from [Theorem 2](#).

We would like to caution against the following intuition. Under SCD preferences, any inadequate-knowledge belief μ has distinct optimal actions at the extreme states of μ 's support. DUB information then guarantees learning because the extreme states can be distinguished from their complements, and so μ is not stationary. While valid for finite states, this is not a generally applicable intuition. Indeed, the following example shows that pairing DUB with distinct optimal actions at all states is not a robust principle for learning.

¹⁸Regardless of the MLRP, DUB implies pairwise distinguishability. To see why the converse is true given the MLRP, consider the case of finite states. Note that for any $\omega' > \omega$, $f(s|\omega')/f(s|\omega) \rightarrow \infty$ as $s \rightarrow \sup S$ (the ratio is increasing by MLRP, and it diverges by pairwise distinguishability); similarly, the ratio goes to 0 as $s \rightarrow \inf S$. Hence, for any ω' and $\varepsilon > 0$, the condition in [Lemma 1](#) is met for $\Omega' = \{\omega'\}$ and $\Omega'' = \{\omega'' : \omega'' < \omega'\}$ when S' is any sufficiently small upper set of signals, while for $\Omega'' = \{\omega'' : \omega'' > \omega'\}$ the condition is met when S' is any sufficiently small lower set. For an infinite state space, the intuition is the same but we appeal to the monotone convergence theorem.

¹⁹“Rich-enough” here means that for any two states, there is a preference in the class such that the optimal actions in those two states are disjoint.

Example 2. 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 $\varepsilon > 0$.²⁰ 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.

Consider normal information. There are full-support priors μ such that the posterior probability $\mu_s(\omega)$ is uniformly bounded away from 1 across signals s and states ω (see [Supplementary Appendix SA.3](#) for details). 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. So [Theorem 1](#) implies inadequate learning. \square

The argument for the necessity of DUB in [Proposition 1](#) is as follows. Take any state ω^* , any two actions $a_1 \neq a_2$, and consider the following SCD utility: for all $\omega < \omega^*$, $u(a_1, \omega) = 1$ and $u(a_2, \omega) = 0$; for all $\omega \geq \omega^*$, $u(a_1, \omega) = 0$ and $u(a_2, \omega) = 1$; and otherwise $u(a, \omega) = -1$. Since all actions except a_1 and a_2 are dominated and can be ignored, [Theorem 2](#) implies that for there to be adequate learning, $\Omega_{a_1, a_2} = \{\omega : \omega < \omega^*\}$ and $\Omega_{a_2, a_1} = \{\omega : \omega \geq \omega^*\}$ must be distinguishable. In particular, ω^* is distinguishable from its lower set. An analogous argument shows that ω^* is distinguishable from its upper set. Since ω^* is arbitrary, DUB holds.

While our main point in this subsection is that DUB is the correct informational condition for adequate learning under SCD preferences, it is also worth noting that for any preferences violating SCD, one can show that there are DUB information structures—e.g., normal information—with inadequate learning at some choice set. In this sense SCD and DUB are a minimal pair of sufficient conditions.

4.2. Learning in a Multi-Dimensional World

We now turn to a multi-dimensional environment: $A, \Omega \subset \mathbb{R}^d$ for some integer $d \geq 1$.²¹ For instance, $A = \{1, 2, 3\}^2$ can represent a set of feasible policies, $\Omega = \{1, 2, 3\}^2$ society’s ideal policy, and individuals have quadratic-loss prefer-

²⁰Strictly speaking, quadratic-loss utility with $\Omega = \mathbb{Z}$ violates our maintained assumption of bounded utility, but we ignore that to keep the example succinct.

²¹We view any $x \in \mathbb{R}^d$ as a column vector and denote its transposition by x' and its standard Euclidean norm by $\|x\|$.

ences $u(a, \omega) = -\|a - \omega\|^2$. Is there a natural class of information structures for which learning obtains?

More generally, consider the following class of preferences:

Definition 5. Preferences are *intermediate* if for all $a_1 \neq a_2$, either $\Omega_{a_1, a_2} = \emptyset$ or $\Omega_{a_1, a_2} = \Omega$ or there are $h \in \mathbb{R}^d$ and $c \in \mathbb{R}$ such that $\Omega_{a_1, a_2} = \{\omega : h \cdot \omega > c\}$.

Introduced by [Grandmont \(1978\)](#), intermediate preferences have preferred sets that are either trivial or half spaces; so if $\omega, \omega' \in \Omega_{a_1, a_2}$, then for any $\lambda \in (0, 1)$ and $\omega'' = \lambda\omega + (1 - \lambda)\omega' \in \Omega$, it holds that $\omega'' \in \Omega_{a_1, a_2}$. A leading family, subsuming quadratic-loss preferences, is weighted Euclidean preferences: $u(a, \omega) = -l((a - \omega)'W(a - \omega))$, for some $d \times d$ symmetric positive definite matrix W and strictly increasing loss function $l : \mathbb{R}_+ \rightarrow \mathbb{R}_+$.²² Another salient example, discussed by [Caplin and Nalebuff \(1988, Section 5\)](#), is the constant-elasticity-of-substitution utility $u(a, \omega) = \left(\sum_{i=1}^d \omega_{(i)}(a_{(i)})^r\right)^{1/r}$ where $a_{(i)}$ and $\omega_{(i)}$ denote the respective i -th coordinates, and $r \neq 0$ is a parameter.

Turning to information, we focus on the familiar class of *location-shift* information structures: $S = \mathbb{R}^d$ and there is a density $g : \mathbb{R}^d \rightarrow \mathbb{R}_{++}$, called the *standard density*, such that $f(s|\omega) = g(s - \omega)$. We restrict attention to standard densities that are uniformly continuous. The following property will be crucial.

Definition 6. A location-shift information structure is *subexponential* if there are $p > 1$ and $M > 0$ such that $g(s) < \exp(-\|s\|^p)$ for all $\|s\| > M$.

A subexponential density has a thin tail in the sense that it eventually decays strictly faster than the exponential density. Our leading example of a subexponential location-shift information structure is multivariate normal information: there is some covariance matrix Σ such that the distribution of signals in state ω is $\mathcal{N}(\omega, \Sigma)$. Here the standard density is that of $\mathcal{N}(0, \Sigma)$, and [Definition 6](#) is verified by taking any exponent $p \in (1, 2)$ and any large $M > 0$. Subexponential information can fail unbounded beliefs; for example, this is the case for normal information when Ω

²²To confirm that these are intermediate preferences, note that by simple algebraic manipulation,

$$(a_1 - \omega)'W(a_1 - \omega) - (a_2 - \omega)'W(a_2 - \omega) = (a_1 - a_2)'W(a_1 + a_2 - 2\omega).$$

Hence, $\omega \in \Omega_{a_1, a_2}$ if and only if $(a_1 - a_2)'W(a_1 + a_2 - 2\omega) > 0$, or equivalently, $h \cdot \omega > c$ where $h = 2(a_2 - a_1)'W$ and $c = (a_2 - a_1)'W(a_1 + a_2)$.

contains non-extreme states, i.e., there is some state in the interior of the convex hull of Ω .

The main result of this subsection is:

Proposition 2. *If preferences are intermediate and the information structure is subexponential location-shift, then there is adequate learning.*

The result follows from [Theorem 2](#) and the next lemma, which says that all half spaces are distinguishable from their complements under subexponential location-shift information. Since the nontrivial preferred sets for intermediate preferences are half spaces, the lemma implies that this combination of preferences and information yields excludability.

Lemma 2. *For a subexponential location-shift information structure, the sets $\{\omega : h \cdot \omega > c\}$ and $\{\omega : h \cdot \omega < c\}$ are distinguishable from each other for any $h \in \mathbb{R}^d$ and $c \in \mathbb{R}$.*

The exponent p being strictly larger than 1 in the definition of subexponential is essential for the lemma. To see that, consider the double-exponential standard density $g(s) = c \cdot \exp(-\|s\|)$ with $c > 0$ a constant of integration. This density is not subexponential, and indeed the conclusion of [Lemma 2](#) fails: no two states $\omega \neq \omega'$ are distinguishable from each other because $f(s|\omega')/f(s|\omega) = g(s - \omega')/g(s - \omega) \leq \exp(\|\omega' - \omega\|)$ for any signal s . The failure of pairwise distinguishability implies inadequate learning even with a binary state when the action set is discrete and preferences are nontrivial.

We can provide an intuition for [Lemma 2](#) by considering a bivariate normal standard density, $g(s) = \exp(-s'\Sigma s/2) / \sqrt{2\pi}$ with Σ a 2×2 covariance matrix. Take an arbitrary hyperplane h , as illustrated in [Figure 3](#). We seek to distinguish the half space to the right of h from its complementary half space to the left. It is sufficient to distinguish an arbitrary single state ω_1 to the right of h from all the states to the left. [Figure 3](#) shows how to construct a sequence of signals verifying that distinguishability. For a sequence of $c_n \rightarrow 0$, select s_n on the iso-density ellipse of level c_n given state ω_1 so that the direction of h is tangent with the ellipse at s_n . For all n , the “ellipsoid distance” between s_n and ω_1 , $\sqrt{(s_n - \omega_1)'\Sigma(s_n - \omega_1)}$, is then smaller than the ellipsoid distance between s_n and any state to the left of h (such as ω_2 and ω_3) by some fixed amount. Due to the normal distribution being subexponential, as $c_n \rightarrow 0$ the likelihood ratio $\frac{g(s_n - \omega)}{g(s_n - \omega_1)} \rightarrow 0$ uniformly across ω to the left of h .

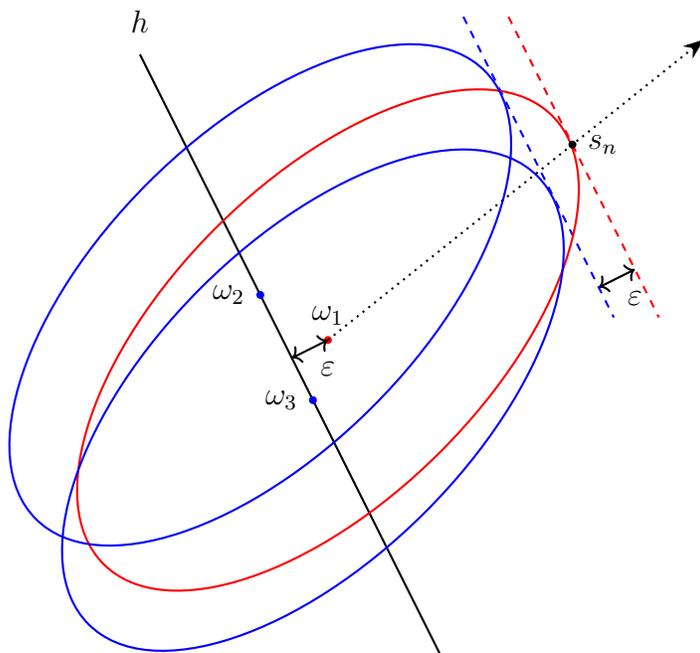


Figure 3: The logic underlying [Lemma 2](#) for a bivariate normal standard density. We seek to distinguish ω_1 from the solid black line. The ellipses are iso-density signals of a given level at the states ω_1, ω_2 , and ω_3 . As s_n grows along the dotted line, corresponding to lower iso-density levels, $\min\{f(s_n|\omega_2)/f(s_n|\omega_1), f(s_n|\omega_3)/f(s_n|\omega_1)\} \rightarrow 0$.

We make two further comments regarding [Proposition 2](#). First, in the one-dimensional environment of [Subsection 4.1](#), SCD is more or less equivalent to preferred sets being half spaces, and DUB is equivalent to the distinguishability of half spaces from their complements in the sense of [Lemma 2](#). [Proposition 2](#) can thus be viewed as an extension of [Proposition 1](#) to a multi-dimensional world; the restriction to location-shift information allows us to unpack the kind of information that yields the requisite half-space distinguishability.

Second, a location-shift information structure does not have to be subexponential to guarantee learning for all intermediate preferences. But it can be shown that if the standard density $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is superexponential in the sense that there are $p \in (0, 1)$ and $M > 0$ such that $g(s) \geq \exp(-\|s\|^p)$ for all $\|s\| > M$, then learning fails for all nontrivial intermediate preferences when A is finite.

5. Theorem 1 and a General Welfare Bound

We now return to the general characterization of adequate learning, [Theorem 1](#), to explain how it is derived. The theorem is best understood as a corollary of a welfare bound regardless of whether there is learning. Stating that result requires some notation. Abusing notation, let

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

be an agent's expected utility when she takes an optimal action under belief μ . Recalling that μ_s denotes the posterior given a belief μ and signal s , let

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

be the expected 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: $\varphi \in \Phi^{BP} \iff \mathbb{E}_{\varphi}[\mu] = \mu_0$. Again abusing notation, we write $u(\varphi) := \mathbb{E}_{\varphi}[u(\mu)]$ for the expected utility of an agent under the distribution of beliefs φ , and analogously write $I(\varphi) := \mathbb{E}_{\varphi}[I(\mu)]$. It follows that

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

is the set of Bayes-plausible distributions of beliefs 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 Bayes-plausible distribution of beliefs is supported on stationary beliefs. Our welfare bound is that eventually all agents are assured a utility level of at least $u_*(\mu_0)$. More precisely:

Theorem 3. *In any equilibrium σ , $\liminf_n \mathbb{E}_{\sigma, \mu_0}[u_n] \geq u_*(\mu_0)$.*

The “if” direction of [Theorem 1](#) readily follows from [Theorem 3](#): when all stationary beliefs have adequate knowledge, a correct action is taken almost surely for any distribution of stationary beliefs, hence $u_*(\mu_0) = u^*(\mu_0)$, and we have adequate learning.

The conclusion of [Theorem 3](#) would be straightforward if we were assured that agents eventually hold stationary beliefs. However, there are networks (with expanding observations) in which with positive probability the beliefs of an infinite number of agents are bounded away from the set of stationary beliefs; see [Example SA.1](#) in [Supplementary Appendix SA.1](#).

Instead, we prove [Theorem 3](#) via an *improvement principle*, as suggested by [Banerjee and Fudenberg \(2004\)](#) and developed by [Acemoglu, Dahleh, Lobel, and Ozdaglar \(2011\)](#) and others. The foundation in our general setting is a novel compactness-continuity argument. First, [Lemma 4](#) in the appendix establishes that Φ^{BP} is compact when both $\Delta\Omega$ and $\Delta\Delta\Omega$ are endowed with the Prohorov metric generated from the metric on Ω . The idea when Ω is countable is that although the prior μ_0 can be supported on an infinite set, it must concentrate an arbitrarily large mass on only finitely many states. Consequently, for any $\delta > 0$, there is a finite subset of states Ω' such that any Bayes-plausible distribution of beliefs puts at least $1 - \delta$ probability on beliefs that put at least $1 - \delta$ probability on Ω' . Using Prohorov’s Theorem, we then deduce that Φ^{BP} is compact. Second, we show that the utility function $u(\varphi)$ and the improvement function $I(\varphi)$ are continuous ([Lemma 5](#) in the appendix), and thus uniformly continuous on 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 distribution of beliefs is in $(\Phi^S)^\varepsilon$, then her ex-ante expected utility is at least close to u_* , as $u(\varphi) \geq u_*$ on Φ^S and $u(\varphi)$ is uniformly continuous. On the other hand, if the 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, hence compact, and $I(\varphi)$ is continuous).

We can then apply an improvement principle. The idea is as follows, where we consider deterministic networks for simplicity. Expanding observations guarantees that we can partition society into “generations” such that an agent in one generation observes a predecessor who is in either the previous generation or the current generation. We inductively argue that the lowest ex-ante utility in each

generation is either close to u_* or increases by a fixed amount compared to the previous generation. Consider an agent’s distribution of social beliefs, φ . Her utility $u(\varphi)$ must be at least the lowest ex-ante expected utility of the previous generation, because the current agent can just mimic the agent with the largest index she observes.²³ Then, as explained in the previous paragraph, either $u(\varphi)$ is at least close to u_* (when φ is in $(\Phi^S)^\varepsilon$), or the agent can improve upon $u(\varphi)$ by at least some fixed amount. 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 3](#).

Although previous authors have deduced versions of [Theorem 1](#) and [Theorem 3](#) 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 to [Theorem 3](#). 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 (an “expert agent”, in their terminology). Our cascade 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; [Remark 5](#) in [Appendix C](#) elaborates. Typically the ranking will be strict, although [Lobel and Sadler \(2015\)](#) note that cascade and diffusion utilities coincide in their binary-state binary-action setting. Methodologically, [Lobel and Sadler’s](#) argument for a minimum improvement, like that of [Acemoglu, Dahleh, Lobel, and Ozdaglar \(2011\)](#), owes to certain monotonicity that does not extend beyond their binary-binary setting.

Remark 3. Our approach to proving [Theorem 3](#) can be adapted to address belief convergence. Since expanding observations is compatible with the observational network having multiple components, one cannot expect the social belief to con-

²³With stochastic networks, the fact that an agent can obtain any observed predecessor’s ex-ante expected utility through mimicking relies on our assumption that players’ observation neighborhoods are drawn independently. Otherwise, whether a player has observed some predecessor may correlate with that predecessor realizing a lower utility.

verge even in probability.²⁴ Furthermore, there can be a positive probability that the social belief is not eventually even in a neighborhood of the set of stationary beliefs, as already noted. Nevertheless, there are reasonable conditions under which convergence to the stationary set does obtain. Consider deterministic networks and assume that society can be covered by finitely many subsequences such that in each subsequence agent n_k observes n_{k-1} . Then, denoting agent n 's (random) social belief by μ_n , it holds that for all $\varepsilon > 0$, $\lim_{n \rightarrow \infty} \Pr(\mu_n \in S^\varepsilon) = 1$, where S^ε denotes the ε -neighborhood of the set of stationary beliefs. See [Proposition SA.1](#) in [Supplementary Appendix SA.1](#). We note that this result applies, in particular, to the immediate-predecessor network and the complete network. The latter is special because the social belief is then a martingale, which is assured to converge almost surely by the martingale convergence theorem. For this case, [Arieli and Mueller-Frank \(2021, Lemma 1\)](#) have established that the limit is stationary.

Remark 4. [Theorem 3](#) can be used to quantify how a failure of excludability impacts welfare. [Proposition SA.2](#) in [Supplementary Appendix SA.2](#) provides a formal result in this vein. In particular, that result implies a sense in which an environment with “approximate excludability” ensures that, eventually, agents’ ex-ante expected utilities are close to the full-information utility.

6. Concluding Remarks

This paper has studied a general model of sequential social learning on observational networks. Our main theme has been how learning turns jointly on preferences and information when there are multiple states. We close by commenting on certain aspects of our approach.

First, our model assumes “non-anonymous sampling”, i.e., whenever an agent sees the action of some predecessor, she knows the identity of that predecessor. However, our methodology extends to anonymous sampling, i.e., when each agent observes only the frequencies of actions in their realized neighborhood, as in [Smith and Sørensen \(2020\)](#). Our results apply in that case when expanding observations (condition (1)) holds for the “induced network structure” $(\tilde{Q}_n)_{n \in \mathbb{N}}$ where each \tilde{Q}_n

²⁴Consider an observational network consisting of two disjoint complete subnetworks: every odd agent observes only all odd predecessors, and symmetrically for even agents. Given any specification in which learning would fail on a complete network—such as the canonical binary state/binary action herding example—there is positive probability of the limit belief among odd agents being different from that among even agents.

is defined by first drawing a neighborhood B_n according to Q_n and then uniformly randomly drawing a single agent from B_n . [Appendix A \(fn. 28\)](#) explains why. Interestingly, the condition coincides with [Smith and Sørensen’s \(2020\)](#) “non-over-sampling” requirement. Note that expanding observations for the induced network (\tilde{Q}_n) is more demanding than expanding observations for (Q_n); this is not surprising since agents have less information when they cannot observe identities. Nevertheless, the requirement is satisfied, for example, when each agent observes the action of a uniform-randomly drawn predecessor or the actions of all predecessors—in either case, not observing their identities. But the requirement is violated when each agent n observes either agent 1 or agent $n - 1$, but doesn’t observe the identity (whereas expanding observations holds here when the identity is observed).

Second, the notion of learning we have adopted considers all possible priors. While this strengthens our sufficiency results, it correspondingly weakens our necessity results. With only two states, learning at any single (nondegenerate) prior is equivalent to learning at all priors. Our earlier working paper ([Kartik, Lee, Liu, and Rappoport, 2022](#), Supplementary Appendix SA.1) provides some analysis concerning the extent to which this is true with multiple states.

Third, our analysis has not touched on the speed of learning/welfare convergence. 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 \(2018\)](#) as well.

Lastly, our work only addresses Bayesian learning with correctly specified agents. 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 \(2020\)](#) and [Bohren and Hauser \(2021\)](#).

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.
- 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*, Providence, Rhode Island: American Mathematical Society.
- BOHREN, J. A. AND D. N. HAUSER (2021): "Learning With Heterogeneous Misspecified Models: Characterization and Robustness," *Econometrica*, 89, 3025–3077.
- CAPLIN, A. AND B. NALEBUFF (1988): "On 64%-Majority Rule," *Econometrica*, 56, 787–814.
- 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. IJIMA, AND Y. ISHII (2020): "Misinterpreting Others and the Fragility of Social Learning," *Econometrica*, 88, 2281–2328.
- 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.
- GRANDMONT, J.-M. (1978): "Intermediate Preferences and the Majority Rule," *Econometrica*, 46, 317–330.
- HANN-CARUTHERS, W., V. V. MARTYNOV, AND O. TAMUZ (2018): "The Speed of Sequential Asymptotic Learning," *Journal of Economic Theory*, 173, 383–409.
- KARTIK, N., S. LEE, T. LIU, AND D. RAPPOPORT (2022): "Beyond Unbounded Beliefs: How Preferences and Information Interplay in Social Learning," Unpublished.
- KARTIK, N., S. LEE, AND D. RAPPOPORT (2023): "Single-Crossing Differences in Convex Environments," *Review of Economic Studies*, forthcoming.
- 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.

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.

Appendices

Appendix A contains the proofs for Theorems 1–3. **Appendix B** contains the proofs of **Proposition 1** and **Lemma 2** (which proves **Proposition 2**). **Appendix C** contains the proof of **Lemma 1** and **Remark 5** on cascade vs. diffusion utility .

A. Backbone Results

In this section, we prove our three theorems in the following setting, which is more general than that described in the main text:

- 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;²⁵
- The utility function $u(a, \omega)$ has absolute value uniformly bounded by \bar{u} and it 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 is given by $Q \equiv (Q_n)_{n \in \mathbb{N}}$, where each Q_n is a probability measure over all neighborhoods, i.e., all subsets of $\{1, 2, \dots, n-1\}$, independent across n , independent of the state ω , and independent of any private signals.

²⁵That is, (Ω, d) is a complete and separable metric space that can be represented as a countable union of compact sets.

When Ω is countable as in the main text, we endow it with the discrete metric so that the sigma-compactness and continuity requirements are trivially satisfied.

Discontinuous utilities. While we make a continuity assumption on preferences, our main results hold for utilities satisfying the following condition that permits discontinuities (cf. [Remark 1](#)):

Condition 1. There is a countable partition of Ω into Borel sets B_i and pointwise equicontinuous functions $v_i : \Omega \rightarrow \mathbb{R}$ uniformly bounded by \bar{u} such that $v_i|_{B_i} = u$.

To obtain our results for such utilities, we can define a new state space $\tilde{\Omega}$ as a disjoint union: $\tilde{\Omega} := \bigsqcup \Omega_i$, where each Ω_i is a copy of Ω . Choose any metric on $\tilde{\Omega}$ that induces the disjoint union topology. Define a utility \tilde{u} on $\tilde{\Omega}$ by $\tilde{u}|_{\Omega_i} := v_i$ for each i . It follows that $\tilde{\Omega}$ is sigma-compact Polish and \tilde{u} is pointwise equicontinuous and uniformly bounded. The information structure is defined such that on each Ω_i it is the same as before. Using our results for the new setting, one can deduce [Theorems 1–3](#) for the original setting.²⁶

A.1. Overarching Probability Space and Beliefs

We now formalize the overarching probability space over all realizations of the state, signals, observation neighborhoods, and actions. We also define formal objects corresponding to agents' social and posterior beliefs and distributions of beliefs.

Overarching probability space. Our probability space is constructed from three components: the Markov kernel F and probability space $(\Omega, \mathcal{B}(\Omega), \mu_0)$; the network structure $Q \equiv (Q_n)_{n \in \mathbb{N}}$; each agent n 's strategy $\sigma_n(\cdot|a_{B_n}, B_n)$ as 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). We suppress the dependence of \mathbb{P} on σ and μ_0 .

²⁶More specifically, the results in the original setting are equivalent to the corresponding results in the new setting restricted to priors/beliefs that put zero probability on the added states $\Omega_i \setminus B_i$. We can use our methodology to derive [Theorems 1–3](#) in the new setting for such restricted beliefs.

Beliefs. Given this overarching probability space, agent n 's social belief (i.e., her belief after observing her neighbors and their actions, but before observing her private signal) is $\mathbb{P}(\cdot|a_{B_n}, B_n)$ and her posterior belief 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. We 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 $\Delta\Omega$) also equipped with the Prohorov metric.

The social 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 social 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 $\Delta\Omega$.

A.2. 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:

$$\int_A \mu(A) d\varphi(\mu) = \mu_0(A), \quad \forall A \in \mathcal{B}(\Omega). \quad (2)$$

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 distribution $\varphi \in \Phi^{BP}$ 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 $\mu_k \rightarrow \mu$ and $\mu_k \in V_\delta$. Since each Ω_i is compact (and thus closed), weak convergence implies

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

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 distribution induced by 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**, we can use Prohorov's theorem again to show:

Lemma 4. Φ^{BP} is compact.

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

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

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

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

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

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

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 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 $\varphi \in \Phi^{BP}$, hence Φ^{BP} is tight. Prohorov's theorem now implies that the closure of Φ^{BP} , which is Φ^{BP} itself, is compact. Q.E.D.

A.3. 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 :

$$u(\mu) := \sup_{a \in A} \int_{\Omega} u(a, \omega) d\mu(\omega).$$

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:

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

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

$$u^*(\mu) := \int_{\Omega} \sup_{a \in A} u(a, \omega) d\mu(\omega).$$

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 use Theorem 2.2.8 in [Bogachev \(2018\)](#), which we restate without proof for our context as the following claim:

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

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

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

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

which converges to 0 as $\mu_k \rightarrow \mu$, by **Claim 1**.

Lastly, suppose we establish that $\Gamma^F := \left\{ \int_S u(\beta(s), \omega) \, dF(s|\omega) \right\}_{\beta: S \rightarrow A}$, as a family of functions of ω indexed by β , is pointwise equicontinuous.²⁷ 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 $\forall \beta$,

$$\begin{aligned} & \left| \int_S u(\beta(s), \omega) \, dF(s|\omega) - \int_S u(\beta(s), \omega') \, dF(s|\omega') \right| \tag{3} \\ & \leq \left| \int_S (u(\beta(s), \omega) - u(\beta(s), \omega')) \, dF(s|\omega) \right| + \left| \int_S u(\beta(s), \omega') \, dF(s|\omega) - \int_S u(\beta(s), \omega') \, dF(s|\omega') \right|. \end{aligned}$$

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 (3) to be smaller than $\varepsilon/2$ (regardless of $\beta(s)$). The second term is smaller than $2\bar{u}d_{TV}(F(\cdot|\omega), F(\cdot|\omega'))$ (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 $d_{TV}(F(\cdot|\omega), F(\cdot|\omega')) < \varepsilon/4\bar{u}$. Therefore, if $d(\omega', \omega) < \min\{\delta_1, \delta_2\}$, then the right-hand side of inequality (3) 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, de-

²⁷Here 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.

fine $u(\varphi) := \mathbb{E}_\varphi[u(\mu)]$, $I(\varphi) := \mathbb{E}_\varphi[I(\mu)]$, 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 $G(\varphi)$.

We note that a belief μ is stationary if and only if $I(\mu) = 0$, and a belief μ has adequate knowledge if and only if $G(\mu) = 0$. To confirm these points, consider stationary beliefs. If there is an action that is a.s. optimal regardless of the signal, then $I(\mu) = 0$. Conversely, if no action 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.

A.4. Proofs for Backbone Results

Logically, **Theorem 3** \implies **Theorem 1** \implies **Theorem 2**. So we prove the results in that order.

Proof of Theorem 3. We prove the result in two steps. In Step 1 below, we prove that if agent n 's social 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 some positive amount bounded away from zero. 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 on Φ^{BP} (as $u(\varphi)$ is continuous, and Φ^{BP} is compact), if $\varphi_n \in (\Phi^S)^\varepsilon$, then $u(\varphi_n) \geq u_* - \gamma(\varepsilon)$ for some $\gamma(\cdot)$ such that $\gamma(\varepsilon) \rightarrow 0$ when $\varepsilon \rightarrow 0$. If, on the other hand, $\varphi_n \in \Phi^+ := \Phi^{BP} \setminus (\Phi^S)^\varepsilon$, then $I(\varphi_n) > 0$ because φ_n puts positive probability on $\{\mu : I(\mu) > 0\}$. Since $(\Phi^S)^\varepsilon$ is open, Φ^+ is a closed subset of a compact set Φ^{BP} ; hence Φ^+ 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 an 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 $\varepsilon \rightarrow 0$ implies $\liminf_n \mathbb{E}_{\sigma, \mu_0}[u_n] \geq u_*$, which completes the proof.

For a given $\varepsilon > 0$, let $\delta = \frac{I(\underline{\varphi})}{4\bar{u}} > 0$, let $N_0 = 1$, and define N_k for $k = 1, 2, \dots$ 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 $n \geq N_k$, $\mathbb{E}_{\sigma, \mu_0}[u_n] \geq \alpha_k := \min\{u_* - \gamma(\varepsilon), \frac{kI(\underline{\varphi})}{2} - \bar{u}\}$. Since $\alpha_0 = -\bar{u}$, 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, $u(\varphi_n) \geq (1-\delta) \cdot \alpha_{k-1} + \delta \cdot (-\bar{u})$.²⁸ 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

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

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 Theorem 1. 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 $\varphi \in \Phi^S$,

²⁸If agents do not observe the identities associated with the observed actions of their predecessors, an agent can uniform-randomly select one of the actions they observe to imitate. So long as the “induced network structure” (i.e., a network structure (\tilde{Q}_n) wherein each \tilde{Q}_n is defined by first drawing a neighborhood B_n from Q_n and then uniform-randomly drawing a single agent from B_n) satisfies expanding observations, the current proof goes through without change using the induced network structure.

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

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

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 3](#) 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 $\mathbb{E}_{\sigma, \mu_0}[u_n] \rightarrow u^*(\mu_0)$. As μ_0 and σ are arbitrarily, we have adequate learning. Q.E.D.

Next we state and prove a more general version of [Theorem 2](#). For any $n \in \mathbb{N}$, define $\Omega_{a_1, a_2}^n := \{\omega : u(a_1, \omega) - u(a_2, \omega) > \frac{1}{n}\}$.

Theorem 2'. *Excludability implies adequate learning at every choice set. There is inadequate learning for choice set $\{a_1, a_2\}$ if Ω_{a_1, a_2} is not distinguishable from Ω_{a_2, a_1}^n for some n .*

Note that when Ω is finite, or the utility difference between any pair of actions is bounded away from zero, a failure of excludability is equivalent to the condition for necessity in the theorem holding for some a_1, a_2 . Hence [Theorem 2](#) is implied by [Theorem 2'](#).

Proof of Theorem 2'. (First statement) First note that excludability (under the full choice set A) implies excludability under any choice subset $A' \subseteq A$. So we fix an arbitrary $A' \subseteq A$ and show that excludability under that subset implies adequate learning at that choice subset. In what follows, the domain of actions should be understood as A' , and we denote a typical element by a' .

[Theorem 1](#) 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, $\mu(\cup_{a'} \Omega_{a', a^*}) > 0$, 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 Ω_{a', a^*}^n 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 Ω_{a', a^*}^n , and hence the open set $\cup_{a'} \Omega_{a', a^*} (= \cup_{a'} \cup_n \Omega_{a', a^*}^n)$, is a union of countably many basic open sets. Since $\mu(\cup_{a'} \Omega_{a', a^*}) > 0$, at least one basic open set contained in Ω_{a', a^*}^n for some a' and n has strictly positive measure, i.e., $\mu(\Omega_{a', a^*}^n) > 0$.

Now denote $\mu'(\cdot) := \mu(\cdot|\Omega_{a^*,a'} \cup \Omega_{a',a^*}^n)$ as the corresponding conditional probability. Since Ω_{a',a^*} is distinguishable from $\Omega_{a^*,a'}$ by excludability, so is Ω_{a',a^*}^n .²⁹ Therefore, for any $\varepsilon > 0$ there exists a set of signals S' such that $\Pr_{\mu'}(S') > 0$ and $\mu'_s(\Omega_{a',a^*}^n) > 1 - \varepsilon$ 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 $(\frac{1}{n}(1 - \varepsilon) - 2\bar{u}\varepsilon)\mu_s(\Omega_{a^*,a'} \cup \Omega_{a',a^*}^n)$, as the expected improvement on $\Omega \setminus (\Omega_{a^*,a'} \cup \Omega_{a',a^*}^n)$ is nonnegative. For small $\varepsilon > 0$, $\frac{1}{n}(1 - \varepsilon) - 2\bar{u}\varepsilon > 0$. Furthermore, integrating $\mu_s(\Omega_{a^*,a'} \cup \Omega_{a',a^*}^n)$ over $s \in S'$ yields $\Pr_{\mu'}(S')\mu(\Omega_{a^*,a'} \cup \Omega_{a',a^*}^n) > 0$. Hence, the ex-ante improvement is bounded below by $(\frac{1}{n}(1 - \varepsilon) - 2\bar{u}\varepsilon)\Pr_{\mu'}(S')\mu(\Omega_{a^*,a'} \cup \Omega_{a',a^*}^n) > 0$. 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 Ω_{a_2,a_1}^n . This means there exists $\mu \in \Delta(\Omega_{a_1,a_2} \cup \Omega_{a_2,a_1}^n)$ with $\mu(\Omega_{a_1,a_2}) > 0$ such that $\mu_s(\Omega_{a_1,a_2}) \leq 1 - \varepsilon$ for some $\varepsilon > 0$ and μ -a.e. s . Consider $\mu' \in \Delta(\Omega_{a_1,a_2} \cup \Omega_{a_2,a_1}^n)$ 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 $\mu'(\cdot|\Omega_{a_2,a_1}^n) = \mu(\cdot|\Omega_{a_2,a_1}^n)$. Under μ' , upon observing signal s , the posterior on Ω_{a_1,a_2} satisfies

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

for μ -a.e. s . Hence, by choosing μ' so that $\frac{\mu'(\Omega_{a_1,a_2})}{\mu'(\Omega_{a_2,a_1}^n)}$ is arbitrarily small, the ratio $\frac{\mu'_s(\Omega_{a_1,a_2})}{\mu'_s(\Omega_{a_2,a_1}^n)}$ 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 $2\bar{u}\mu'_s(\Omega_{a_1,a_2}) - \frac{1}{n}\mu'_s(\Omega_{a_2,a_1}^n)$, which is strictly negative when $\frac{\mu'_s(\Omega_{a_1,a_2})}{\mu'_s(\Omega_{a_2,a_1}^n)}$ 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. **Theorem 1** implies there is inadequate learning for choice set $\{a_1, a_2\}$. Q.E.D.

B. Applications

We now specialize to the main text's setting: Ω is countable, endowed with the discrete metric, and $F(\cdot|\omega)$ are absolutely continuous with respect to each other, and so there are densities $f(\cdot|\omega) > 0$.

²⁹ In fact, excludability is equivalent to: Ω_{a_1,a_2}^n is distinguishable from Ω_{a_2,a_1} for all a_1, a_2 and n .

B.1. SCD Preferences & DUB Information

Proof of Proposition 1. Sufficiency follows directly from [Theorem 2](#). For necessity, first observe that if the information structure fails DUB, then there exists some state ω^* such that ω^* is not distinguishable from its lower set (or from its upper set, which has a symmetric argument). Fix any pair of distinct actions a_1 and a_2 , and define the following SCD preferences: for $\omega < \omega^*$, $u(a_1, \omega) = 1$ and $u(a_2, \omega) = 0$; for $\omega \geq \omega^*$, $u(a_1, \omega) = 0$ and $u(a_2, \omega) = 1$; and any other actions are strictly dominated. It follows that Ω_{a_2, a_1} is not distinguishable from $\{\omega : u(a_1, \omega) - u(a_2, \omega) > \frac{1}{2}\}$. By [Theorem 2'](#), there is inadequate learning when the choice is $\{a_1, a_2\}$, and since all other actions are strictly dominated, also for the full choice set A . Q.E.D.

B.2. Intermediate Preferences & Location-Shift Information

The proof of [Lemma 2](#) is more involved than the intuition given in the main text using [Figure 3](#), because in general one cannot explicitly identify the sequence of signals that establishes distinguishability of the relevant two sets.

We will use the following claim in proving [Lemma 2](#). For any $h, x \in \mathbb{R}^d$, let $\|x\|_h := h \cdot x$ be the “signed distance” of x in direction h , i.e., between x and the hyperplane $\{z : h \cdot z = 0\}$. Note that $\|\cdot\|_h$ is linear, so $\|x - x'\|_h = \|x\|_h - \|x'\|_h$.

Claim 2. *If a standard density g is subexponential, then for any \bar{s} with $\|\bar{s}\|_h > 0$, and $\varepsilon \in (0, 1)$, there is s with $\|s - \bar{s}\|_h \geq 1$ such that:*

1. $\sup_{\{s' : \|s' - s\|_h \geq 1/\|\bar{s}\|_h\}} \frac{g(s')}{g(s)} < \varepsilon$; and
2. $\sup_{\{s' : 0 < \|s' - s\|_h < 1/\|\bar{s}\|_h\}} \frac{g(s')}{g(s)} < 2$.

Proof. Suppose not, to contradiction. Then there exists \bar{s} with $\|\bar{s}\|_h > 0$ and $\varepsilon \in (0, 1)$ with the following property: for every s with $\|s - \bar{s}\|_h \geq 1$, we can find s' with $\|s' - s\|_h > 0$ such that either (i) $\|s' - s\|_h \geq 1/\|\bar{s}\|_h$ and $\frac{g(s')}{g(s)} \geq \varepsilon$, or (ii) $0 < \|s' - s\|_h < 1/\|\bar{s}\|_h$ and $\frac{g(s')}{g(s)} \geq 2$. For an arbitrary choice of s' given s , we define $k_s := \|s' - s\|_h$. That means, for each s with $\|s - \bar{s}\|_h \geq 1$, we have $k_s > 0$ and a signal s' with $\|s' - s\|_h = k_s$ such that either (i) $\frac{g(s')}{g(s)} \geq \varepsilon \geq \varepsilon^{k_s \|\bar{s}\|_h}$ (because $k_s \|\bar{s}\|_h \geq 1$), or (ii) $\frac{g(s')}{g(s)} \geq 2 > \varepsilon^{k_s \|\bar{s}\|_h}$ (because $\varepsilon < 1$).

We construct a sequence of signals $(s_i)_{i=1}^\infty$. First, take any s_1 such that $\|s_1 - \bar{s}\|_h = 1$. Then, for all $i > 1$, take any s_i given s_{i-1} as explained in the previous paragraph. Note that for all i , $\|s_i - s_{i-1}\|_h = k_{s_{i-1}}$, so $\|s_i\|_h = (\|\bar{s}\|_h + 1) + \sum_{j=1}^{i-1} k_{s_j}$.

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

$$(\|s_i\|_h - \|\bar{s}\|_h - 1) \|\bar{s}\|_h \log(\varepsilon) + \log(g(s_1)) \leq \log(g(s_i)). \quad (4)$$

However, since g is subexponential, and $\|s_i\|_h \leq \|s_i\| \|h\|$, there is $p > 1$ such that for all large enough i ,

$$\log(g(s_i)) < - \left(\frac{\|s_i\|_h}{\|h\|} \right)^p. \quad (5)$$

The left-hand side of inequality (4) is linear in $\|s_i\|_h$ while the right-hand side of inequality (5) has exponent $p > 1$, so for large enough i these inequalities are in contradiction.

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

Proof of Lemma 2. Without loss, we only prove that $\{\omega : h \cdot \omega > c\}$ is distinguishable from $\{\omega : h \cdot \omega < c\}$.

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

As noted after **Definition 1**, it is sufficient to prove that any $\bar{\omega} \in \{\omega : h \cdot \omega > c\}$ is distinguishable from $\{\omega : h \cdot \omega < c\}$.³⁰ So take any such $\bar{\omega}$ and μ with $\mu(\bar{\omega}) > 0$. Define $\bar{s}_i := s_i^* + \bar{\omega}$. It follows that for all i ,

$$\|\omega - \bar{\omega}\|_h < 0 \implies \frac{f(\bar{s}_i|\omega)}{f(\bar{s}_i|\bar{\omega})} = \frac{g(\bar{s}_i - \omega)}{g(\bar{s}_i - \bar{\omega})} = \frac{g(s_i^* + (\bar{\omega} - \omega))}{g(s_i^*)} < 2, \quad (6)$$

and

$$\|\omega - \bar{\omega}\|_h \leq -\frac{1}{\|s_{i-1}^*\|_h} \implies \frac{f(\bar{s}_i|\omega)}{f(\bar{s}_i|\bar{\omega})} = \frac{g(s_i^* + (\bar{\omega} - \omega))}{g(s_i^*)} < \frac{1}{i-1}, \quad (7)$$

³⁰We note that this uses the assumption of countable states.

and thus,

$$\begin{aligned} \frac{\mu(\{\omega : h \cdot \omega < c\}|\bar{s}_i)}{\mu(\bar{\omega}|\bar{s}_i)} &\leq \frac{\sum_{\|\omega-\bar{\omega}\|_h < 0} \mu(\omega) f(\bar{s}_i|\omega)}{\mu(\bar{\omega}) f(\bar{s}_i|\bar{\omega})} \\ &< \frac{1}{i-1} \frac{\sum_{\|\omega-\bar{\omega}\|_h \leq -1/\|s_{i-1}^*\|_h} \mu(\omega)}{\mu(\bar{\omega})} + 2 \frac{\sum_{-1/\|s_{i-1}^*\|_h < \|\omega-\bar{\omega}\|_h < 0} \mu(\omega)}{\mu(\bar{\omega})}. \end{aligned} \quad (8)$$

The last expression can be taken arbitrarily small because $\|s_{i-1}^*\|_h \rightarrow \infty$ as $i \rightarrow \infty$.

It remains only to show that the above argument holds for a positive measure of signals rather than just a single \bar{s}_i . Since g is uniformly continuous, there is a neighborhood of \bar{s}_i , say \bar{S}_i , over which (6) and (7) hold with slightly relaxed bounds; for instance, the bounds can be relaxed to 4 and $2/(i-1)$, respectively. This establishes the analog of inequality (8) for all signals in \bar{S}_i with the relaxed bounds. Since we have assumed $g(\cdot) > 0$, each \bar{S}_i has positive measure, so we conclude $\bar{\omega}$ is distinguishable from $\{\omega : h \cdot \omega < c\}$. *Q.E.D.*

C. Other Material

Proof of Lemma 1. As noted before the lemma, Ω' is distinguishable from Ω'' if and only if each $\omega' \in \Omega'$ is distinguishable from Ω'' . So fix any $\omega' \in \Omega'$.

We first prove that if the lemma's condition holds, then ω' is distinguishable from Ω'' . Take any probability measure $\mu \in \Delta(\{\omega'\} \cup \Omega'')$ such that $\mu(\omega') > 0$. By assumption, for any $\varepsilon > 0$ there exists a positive-probability set of signals S' such that $\frac{f(s|\omega')}{f(s|\omega'')} < \varepsilon, \forall \omega'' \in \Omega'', \forall s \in S'$. It follows that for all $s \in S'$,

$$\mu(\omega'|s) = \frac{f(s|\omega')\mu(\omega')}{\sum_{\tilde{\omega} \in \{\omega'\} \cup \Omega''} f(s|\tilde{\omega})\mu(\tilde{\omega})} = \frac{\mu(\omega')}{\mu(\omega') + \sum_{\tilde{\omega} \in \Omega''} \frac{f(s|\tilde{\omega})}{f(s|\omega')} \mu(\tilde{\omega})} > \frac{\mu(\omega')}{\mu(\omega') + \varepsilon}.$$

Since for any $\varepsilon > 0$ we can find a positive-probability set of signals S' satisfying the above inequality, we conclude that for any $\varepsilon > 0$, $\Pr_\mu(s : \mu_s(\Omega') > 1 - \varepsilon) > 0$.

We next prove that if ω' is distinguishable from Ω'' , and Ω'' is finite, then the lemma's condition holds. Consider any μ uniformly distributed over $\{\omega'\} \cup \Omega''$. The distinguishability of ω' from Ω'' implies that for every $\varepsilon > 0$ there is a positive-probability set of signals S' such that $\forall s \in S'$ we have $\frac{\sum_{\tilde{\omega} \in \Omega''} f(s|\tilde{\omega})}{f(s|\omega')} < \varepsilon$, and so $\frac{f(s|\tilde{\omega})}{f(s|\omega')} < \varepsilon$ for every $\tilde{\omega} \in \Omega''$. *Q.E.D.*

Remark 5. 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 distribution of beliefs that is supported on the set of feasible posteriors (i.e., those available under the given information structure and the prior); call the corresponding signal structure the expert signal structure.

Let us now argue that diffusion utility is lower than cascade utility. Notice that diffusion utility must be lower than first drawing a posterior from an arbitrary Bayes-plausible distribution of stationary beliefs and then drawing a signal from the expert signal structure, because this “combined” signal structure is Blackwell more informative than just the expert signal structure. But in the combined structure, the expert signal has no value by definition of stationary beliefs, and so the combined signal structure provides a utility equal to that from the (arbitrary Bayes-plausible) distribution of stationary belief distributions.

Diffusion utility and cascade utility coincide with two states and two actions, as noted by Lobel and Sadler’s (2015). But adding even one action can break this coincidence, e.g., if the third action is a “safe” action—one that is optimal only for some interval of interior beliefs—that shrinks the set of stationary beliefs. For a starker example, recall Example 1 with $\Omega = \{0, 1\}$, $A = [0, 1]$, and $u(a, \omega) = -(a - \omega)^2$. Any nontrivial information structure leads to learning, 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.