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

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

December 25, 2022

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 sufficient—and, in a sense, necessary—condition for learning dubbed *excludability*. Excludability is a joint property of agents' preferences and their information. When required to hold for all preferences, it is equivalent to information having "unbounded beliefs", which demands that any agent can individually identify the truth, even if only with small probability. But unbounded beliefs may be untenable with more than two states: e.g., it is incompatible with the monotone likelihood ratio property. Excludability reveals that what is crucial for learning, instead, is that a single agent must be able to rule out any wrong action, even if she cannot take the correct action. Consequently, excludability helps study classes of preferences and information that mutually ensure learning. We develop two such pairs: (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, Euclidean preferences and subexponential location-shift information.

Keywords: social learning; herds; information cascades; single crossing; Euclidean preferences; unbounded beliefs.

---

*We thank Nageeb Ali, Marina Halac, Ben Golub, Andreas Kleiner, Elliot Lipnowski, José Montiel Olea, Xiaosheng Mu, Harry Pei, Evan Sadler, Lones Smith, and Peter Sørensen for helpful comments. We also received useful feedback from various seminar and conference audiences. César Barilla, John Cremin, and Zikai Xu provided excellent research assistance. Kartik gratefully acknowledges support from NSF Grant SES-2018948.

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

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

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

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

1. Introduction

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

One would anticipate that whether there is social learning depends on the combination of agents' preferences and their information structure. But economists have largely emphasized, at least when the action set is finite, 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 (non-trivial) 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.

Our work is rooted in the observation that with multiple—i.e., more than two—states, it is essential to understand the interplay of preferences and information for social learning. Even though unbounded beliefs still characterizes learning for all preferences (Arieli and Mueller-Frank, 2021),[2] it is now an unduly 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 normal distributions 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 how certain one can become about any non-

[1] Unless noted otherwise, our introduction should be understood as referring to the canonical sequential social learning model with a finite action space, homogeneous preferences, and no direct payoff externalities. It is well understood 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.

[2] Arieli 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).

extremal state based on observing one signal.[3] Is social learning doomed with multiple states for familiar information structures like normal information?

Excludability. Our paper shows that the answer is no. With multiple states, unbounded beliefs is not necessary for learning with standard preferences. In general, whether there is learning depends jointly on preferences and information. We present a simple joint condition, *excludability*, and show that it 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.[4]

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. An intuition for why excludability is sufficient for learning is that society can never get stuck on a wrong action: if an action a' is suboptimal at the true state, then because the correct action, say a^*, is better than a', some agent will receive a private signal convincing her not to take action a'. Notably, it may be impossible for the agent's signal to actually induce her to take the correct action a^*—but because her private information convinces her that a^* is better than a', she must take some action other than a'. (We give a concrete example later in the introduction.) Excludability turns out to guarantee that the process of social learning eventually leads agents to settle on the action a^*. By contrast, under unbounded beliefs, a single agent's private signal can lead her to immediately take the correct action no matter her current (full-support) belief. We view the distinction of social learning arising from the individual capacity to *rule out wrong actions* rather than to discover the correct action as an insight of our paper; this distinction cannot be seen with only two states.

Excludability provides a useful perspective on existing ideas in the social learning literature. For instance, as detailed in Section 3, an information structure yields excludability for

[3] So binary states is special because all states are extreme states; in general, there need not even be any extreme states (e.g., Ω is the set of integers). There is nothing special about normal information violating unbounded beliefs. Under full support (i.e., no signal excludes any state), any information structure satisfying the widely-used monotone likelihood ratio property precludes unbounded beliefs with multiple states; see Remark 1 in Section 4.

[4] The sense in which is it necessary entails (i) varying the choice set that society may face, and (ii) a technical qualifier that is satisfied, for example, when the state space is finite. We also note that although we say "learning" in the introduction for brevity, we actually study adequate learning à la Aghion, Bolton, Harris, and Jullien (1991). The distinction is not crucial, and it is moot when there is a distinct optimal action at every state.

all preferences if and only if that information structure has unbounded beliefs; moreover, when there are only two states, unbounded beliefs is equivalent to excludability under any non-trivial preference. On the other hand, with any number of states, any information structure with bounded beliefs fails excludability when paired with any non-trivial preferences.

More importantly, we can use excludability to deduce tenable conditions on information—weaker than unbounded beliefs—that are sufficient for social learning in general observational networks for canonical classes of preferences. This approach to relaxing stringent conditions, informational or otherwise, is classical in other areas of economics,[5] but is novel to social learning.

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 economics: it captures settings in which there are no preference reversals as the state increases; in particular, it is satisfied by any supermodular utility function. 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—in fact, in a sense minimally sufficient—for excludability, and hence learning. For a direct intuition on how DUB information and SCD preferences interplay, consider normal information again. There are preferences for which learning can fail because society may get stuck at some belief at which agents are taking an incorrect action, but only a strong signal about an intermediate state can change the action—alas no such signal is available. However, under SCD preferences, if a strong signal about some intermediate state ω would

[5] For instance, first-order stochastic dominance is refined to second-order stochastic dominance by restricting to concave (and increasing) utility functions. Another exemplar is refining the Blackwell (1953) "sufficiency" partial order on information structures to "accuracy" by restricting preferences to satisfy single crossing or interval dominance (Lehmann, 1988; Persico, 2000; Quah and Strulovici, 2009).

change the action, then so would either a signal that indicates the state is very likely to be at least ω or at most ω. Such signals are guaranteed in the normal example, and more generally by DUB information.

To illustrate concretely how learning obtains without unbounded beliefs, consider normal information with three states, $\Omega = \{1, 2, 3\}$, and agents choosing an action $a \in A = \{1, 2, 3\}$ with the SCD preferences depicted in Figure 1. The correct action in each state ω is $a = \omega$. When each agent observes all predecessors' actions, Figure 1 shows two numerically-simulated paths of "public beliefs" given the true state $\omega = 2$.[6] The public 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. Notice that—consistent with a failure of unbounded beliefs—there is a range of beliefs, indicated by the grey shading, such that for any public 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, though, after enough switching between actions 1 and 3, the public belief is driven outside the grey region and it becomes possible for an agent to actually take the correct action 2. Eventually, society settles on that action.

Euclidean preferences. Our second application in Section 4 is to *weighted Euclidean preferences* in multidimensional spaces. Agents' preferences over actions $a \in \mathbb{R}^d$ in state $\omega \in \mathbb{R}^d$ are represented by the utility function $u(a, \omega) = -l\left((a - \omega)' W (a - \omega)\right)$, where W is a $d \times d$ symmetric positive definite matrix and l is a strictly increasing loss function.[7] The canonical Euclidean preferences are the special case $u(a, \omega) = -\sum_{i=1}^{d}(a_i - \omega_i)^2$. Such utility functions are commonly assumed to model "spatial preferences", including in political economy and communication/delegation; they capture the notion that the state represents the ideal action or bliss point, and actions are less preferred the further they are in (weighted) Euclidean distance from the bliss point.

Using excludability, we show that social learning obtains under weighted Euclidean preferences so long as information is given by a *subexponential location-shift family*. Location-shift families are widely-used information structures: for some density $g : \mathbb{R}^d \to \mathbb{R}^d$, the

[6] The public belief at time n is agent n's belief about the state given only the history of predecessors' actions, prior to observing her own private signal.

[7] Here a and ω are viewed as column vectors and $'$ denotes transposition.

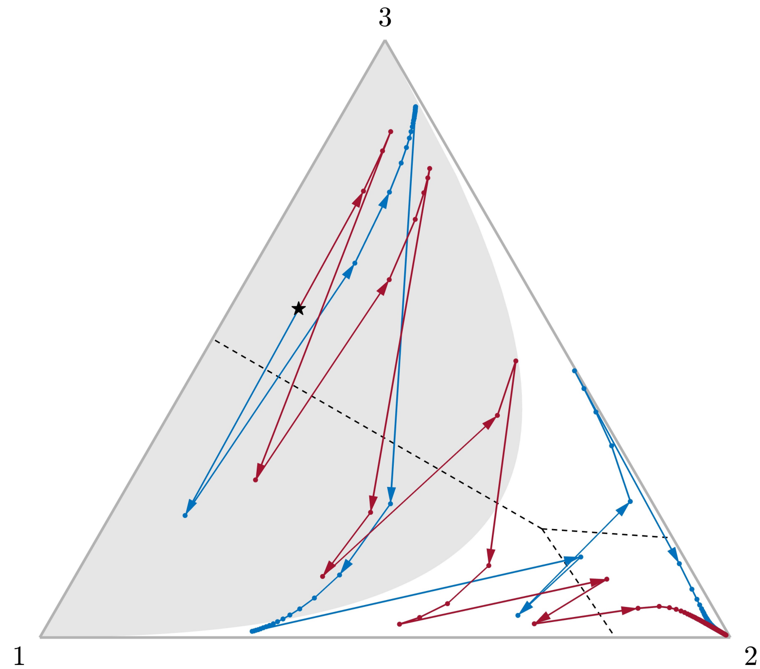

Figure 1: Two simulated public belief paths—one in red and one in blue—in a complete network. There are three states labeled $1, 2, 3$, and there is normal information (with standard deviation 1.2). There are three actions with respective state-contingent utilities $(1, 0, -0.3)$, $(0, 0.2, 0)$, and $(-0.3, 0, 1)$. So SCD holds; 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 signal can lead to state 2's correct action. On each path, a dot represents the public belief after an agent has acted, and arrows indicate the sequencing.

signal distribution in any state θ is given by $g(s - \omega)$. Loosely, our 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 guarantees excludability under weighted Euclidean preferences. 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. Our paper also makes a methodological contribution in developing 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 it is optimal for an agent to ignore the signal she receives from the given information structure, whereas an adequate-knowledge belief is one at which there is an action that is optimal across all states in the belief's support. Excludability is a simple sufficient condition for all stationary beliefs to have adequate knowledge; subject to a technical qualifier, it is also necessary when we require that learning must obtain no matter which choice set society may face (which, as we explain, may be unavoidable when one seeks learning for a broad class of preferences).

Theorem 1 itself is a consequence of Theorem 3 in Section 5, which is a welfare result of independent interest because it applies even when learning fails. This theorem establishes a welfare lower bound: roughly, 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 yields 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 for binary states, Smith and Sørensen (2000) introduce the distinction between bounded and unbounded beliefs; they show that, given any non-trivial preference, there is learning if and only if beliefs are unbounded. For the complete network but with multiple states, Arieli and Mueller-Frank (2021, Theorem 1) show that unbounded beliefs—which they call "totally unbounded beliefs"—is sufficient for learning, and also necessary if learning must obtain no matter society's preferences.[8] Importantly, the approach of both Smith and Sørensen (2000) and Arieli and Mueller-Frank (2021) rests on the fact that in the complete network, the public belief (see fn. 6) is a martingale.

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. They only allow for binary states and binary actions, however. They introduce the condition of expanding observations on the observational network structure, showing that it is necessary for learning. Their Theorem 2 establishes that it is also sufficient for learning with unbounded beliefs. Building on Banerjee and Fudenberg (2004), a key contribution of Acemoglu, Dahleh, Lobel, and

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

Ozdaglar (2011) is to use an *improvement principle* to deduce learning (under unbounded beliefs and expanding observations); this approach works even though martingale arguments fail. Lobel and Sadler (2015) introduce a notion of "information diffusion" and use the improvement principle to establish information diffusion even when learning fails.

The analysis in both Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) and Lobel and Sadler (2015) relies on their binary-state binary-action structure.[9] 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. We use that to establish the cascade-utility welfare bound (and, consequently, our learning characterizations) in any network satisfying expanding observations.

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

We leave to future research the speed of learning/welfare convergence. 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, we note that there is a large literature on non-Bayesian social learning, surveyed by Golub and Sadler (2016). There has also been recent interest in (mis)learning among misspecified Bayesian agents; see, for example, Frick, Iijima, and Ishii (2020, 2021) and Bohren and Hauser (2021).

2. Model

In the main text of the paper, we focus on a model that simplifies some technical issues. Our Theorems 1, 2, and 3 hold in a more general setting described in Appendix A.

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

There is a countable (finite or infinite) set of states $\Omega \subset \mathbb{R}^d$, and standard Borel spaces of actions A and signals S. To avoid trivialities, we assume $|A| > 1$. An information or signal structure is given by a collection of probability measures over S, one for each state, denoted by $F(\cdot|\omega)$. Assume that for any ω and ω', $F(\cdot|\omega)$ and $F(\cdot|\omega')$ are mutually absolutely continuous. It follows that each $F(\cdot|\omega)$ has a density $f(\cdot|\omega)$; more precisely, this is the Radon-Nikodym derivative of $F(\cdot|\omega)$ with respect some reference measure that is mutually absolutely continuous with every $F(\cdot|\omega')$. Without further loss of generality we assume $f(s|\omega) > 0$, so that no signal excludes any state.

The game. At the outset, date 0, a state ω is drawn from a common prior probability mass function $\mu_0 \in \Delta\Omega$. (For any measurable space X, ΔX denotes the set of probability measures over X.) An infinite sequence of agents, indexed by $n = 1, 2, \ldots$, then sequentially select actions. An agent n observes both a private signal s_n drawn independently from $f(\cdot|\omega)$ and the actions of some subset of her predecessors $B_n \subseteq \{1, 2, \ldots, n-1\}$, and then chooses her action $a_n \in A$. No agent observes either the state or any of her predecessors' signals. Each observational neighborhood B_n is stochastically generated according to a probability distribution Q_n over all subsets of $\{1, 2, \ldots, n-1\}$, assumed to be independent across n, independent of the state ω, and independent of any private signals. The distributions $(Q_n)_{n\in\mathbb{N}}$ constitute the observational network structure and are common knowledge, but the realization of each neighborhood B_n is the private information of agent n.[10]

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

All agents are expected utility maximizers and share the same preferences over their actions: every agent n's preferences are represented by the utility function $u : A \times \Omega \to \mathbb{R}$. To ensure expected utility is well defined for all probability distributions, we assume that utility is bounded: there is $\overline{u} \geq 0$ such that $|u(\cdot, \cdot)| \leq \overline{u}$.

We study the Bayes Nash equilibria—hereafter simply equilibria—of this game. An equilibrium exists if for every belief there is an optimal action, which we will implicitly assume to be the case.[11]

[10] We make this assumption for concreteness; our analysis applies even if agents receive arbitrary information about their predecessors' realized neighborhoods.

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

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

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

Given a prior μ_0 and a strategy profile σ, agent n's utility u_n is a random variable. Let $\mathbb{E}_{\sigma,\mu_0}[u_n]$ be agent n's ex-ante expected utility. We say there is *adequate learning* if for every prior μ_0 and every equilibrium strategy profile σ, $\mathbb{E}_{\sigma,\mu_0}[u_n] \to u^*(\mu_0)$. In words, adequate learning requires that given any prior and equilibrium, no matter which state is realized, eventually agents take actions that are arbitrarily close to optimal in that state. We say there is *inadequate learning* if adequate learning fails.[12]

We will also be interested in whether there is (in)adequate learning when agents have to choose from a subset of actions, referred to as a *choice set*. Variation in choice sets is familiar from the comparative-statics literature, and in a social learning context, it can be subsumed when one is interested in learning for a rich-enough class of preferences—specifically, when any subset of actions may be dominated. Accordingly, we say that there is *(in)adequate learning for a choice set* $\tilde{A} \subseteq A$ if there is inadequate learning when agents are restricted to choose from actions in $\tilde{A}$.

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

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

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

Accordingly, we assume expanding observations. Here are a few leading examples of network structures with expanding observations: (i) the classic complete network in which each agent's neighborhood is all her predecessors (formally, $Q_n(B_n = \{1, \ldots, n-1\}) = 1$);

[12] That 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. 11). 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.

(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 $k \in \{1, \ldots, n-1\}$).

Anonymous sampling. Our model assumes that each agent observes the identity of the agent taking each action in the (sub-)history she observes, i.e., in which period each observed action was taken. Our analysis extends to various cases of "random sampling" in which predecessors' identities are not observed. We elaborate in the Conclusion.

3. Characterizations of Learning

3.1. Stationary Beliefs and Adequate Knowledge

The backbone for our analysis 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.

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

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 elaborate on Theorem 1 in Section 5, instead turning now to how we build on it for a more practicable characterization of learning.

3.2. Excludability

Theorem 1's condition of all stationary beliefs having adequate knowledge is not transparent about the kind of information that generates adequate learning for any given preferences, and it may also be difficult to verify when the state/signal spaces are large. Accordingly, it is of interest to identify a simpler condition on primitives.

To that end, a key notion is whether information allows an agent to become arbitrarily sure about a subset of states Ω' relative to another (disjoint) subset Ω'', no matter what belief the agent starts from that does not preclude Ω'. Formally, writing $\mu(\Omega'|s)$ for the posterior on Ω' generated from belief μ and signal s:

Definition 1. A set Ω' is *distinguishable* from another set Ω'' if for any $\mu \in \Delta(\Omega' \cup \Omega'')$ with $\mu(\Omega') > 0$, it holds that $\operatorname{Supp} \mu(\Omega'|\cdot) \ni 1$.

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 in terms of 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.

It bears emphasis 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*. To wit, in Figure 2 each ω' is distinguishable from each $\omega'' \neq \omega'$, but no ω' is distinguishable from $\Omega \setminus \{\omega'\}$.[13] For instance, looking at $\omega' = 3$ in the figure, a signal leading to a posterior close to the 1–3 edge (resp., the 2–3 edge) distinguishes state 3 from 2 (resp., from 1), but there is no signal leading to a posterior close to the 3 vertex.

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

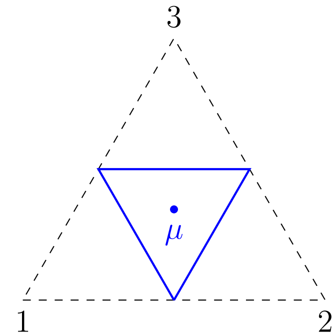

Figure 2: $\Omega = \{1, 2, 3\}$, and the set of posteriors under prior μ is indicated by the solid triangle. So each state is distinguishable from every other, but no state is distinguishable from its complement.

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 unbounded beliefs is highly restrictive with more than two states; for example, with ordered states and signals, it is incompatible with the ubiquitous monotone likelihood ratio property (MLRP) property in information economics:[14]

Remark 1. Under any MLRP information structure, no state ω is distinguishable from any $\{\omega', \omega''\}$ if $\omega' < \omega < \omega''$. In particular, MLRP is incompatible with unbounded beliefs.

Fortunately, learning only requires certain subsets of states to be distinguished from each other. Given utility function $u(a, \omega)$ and any two actions a_1 and a_2, let the preferred set $\Omega_{a_1,a_2} := \{\omega : \ u(a_1, \omega) - u(a_2, \omega) > 0\}$ 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 pair of actions a_1 and a_2, Ω_{a_1,a_2} is distinguishable from Ω_{a_2,a_1}.

Excludability is a simple joint condition on preferences and information. In words, 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 that event. Plainly, given any non-trivial preferences, excludability fails if there is *bounded beliefs*: no state is distinguishable from any other.[15] On the other hand, an information structure yields excludability for all preferences if and only if it has unbounded beliefs: sufficiency is immediate; for necessity, note that if ω is not distinguishable from its complement,

[14] For ordered state and signals spaces, a signal density $f(s|\omega)$ satisfies the MLRP if $\forall s' > s$ and $\forall \omega' > \omega$, $f(s|\omega')/f(s|\omega) \leq f(s'|\omega')/f(s'|\omega)$.

[15] Preferences are non-trivial if for some two actions and two states, one action is strictly preferred at one state and the other action at the other state.

then excludability fails when preferences are such that for some a_1 and a_2, $\Omega_{a_1,a_2} = \{\omega\}$ while $\Omega_{a_2,a_1} = \Omega \setminus \{\omega\}$. Furthermore, with only two states, excludability under any non-trivial preferences is equivalent to unbounded beliefs. The crucial point, however, as we will delve into in Section 4, is that with multiple states, interesting combinations of preferences and information yield excludability even absent unbounded beliefs. 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.*

Theorem 2 follows from Theorem 1, but excludability is easier to interpret and often more tractable than the condition in Theorem 1. Indeed, excludability is sufficient for adequate learning at every choice set because excludability guarantees that, at every choice set, every stationary belief has adequate knowledge. To see why, suppose a belief μ has *in*adequate knowledge, so that $c(\mu) \neq c(\omega^*)$ for some state $\omega^* \in \operatorname{Supp}\mu$. (For simplicity, assume $c(\mu)$ and $c(\omega^*)$ are singletons.) Excludability implies that state ω^* is distinguishable from the set of states $\Omega_{c(\mu),c(\omega^*)}$. Hence, with positive probability, an agent who starts with belief μ will obtain a posterior that puts arbitrarily large probability on ω^* 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 in the previous paragraph's argument, excludability only ensures that a single agent can rule out any wrong action—say, $c(\mu)$ when the true state is ω^*—in favor of the correct action $c(\omega^*)$.[16] But it could be that under belief μ, the agent never takes the correct action, no matter her signal. Recall the example discussed in the introduction, illustrated in Figure 1, where it is impossible for the first few agents to take the correct action when the state is 2. By contrast, social learning has traditionally focused on a different mechanism: it is always possible for an agent to take the correct action. That has lead to the emphasis on unbounded beliefs. When there are two states and finite actions, always being able to rule out a wrong action and always being able to take the correct action are equivalent, as they both simply reduce to unbounded beliefs. But not so more generally, as underlined by the sufficiency of excludability for adequate learning.

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 that when the choice set is $\tilde{A} = \{a_1, a_2\}$, a

[16] To put an even finer point on it, all that is needed for adequate learning is that every wrong action can be ruled out in favor of *some* action, not necessarily the correct action. Under excludability, though, wrong actions can be ruled out by the correct action.

belief that puts small probability on Ω_{a_1,a_2} relative to Ω_{a_2,a_1} is stationary and has inadequate knowledge.[17] Hence, Theorem 1 implies that excludability is necessary for learning when we seek learning for all choice sets. The following example illustrates.

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 non-trivial 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 the choice set A. But absent excludability there is *in*adequate 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, the class must be such that for any pair of actions, there is a preference in the class such that only these two actions are undominated. For a simple illustration, notice that the class of all preferences is obviously broad enough in this sense, and an information structure yields excludability for all preferences if and only if it has unbounded beliefs (because, given any state ω, preferences may only depend on whether the state is ω). Hence, Theorem 2 immediately implies:

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

Although it is not conceptually novel, we state the above result because it has not been established previously for general observational networks with multiple states. Arieli and Mueller-Frank (2021, Theorem 1) have established the result for a complete network and Acemoglu, Dahleh, Lobel, and Ozdaglar (2011, Theorem 2) for general networks but with only two states.

As mentioned in the introduction, with a discrete action space scholars have focused on unbounded beliefs as the key condition for learning. The only exception we are aware of is

[17] The qualifier that Ω is finite is technical; it ensures that $\min_{\omega\in\Omega_{a_1,a_2}}[u(a_1,\omega)-u(a_2,\omega)]$ exists. The appendix provides a stronger version of Theorem 2, labelled Theorem 2′, that drops finiteness.

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. We note that this too follows from Theorem 2; indeed, the theorem implies that learning obtains for general observational networks. For, under simple utility, given any two actions a_1 and a_2, the preferred sets Ω_{a_1,a_2} and Ω_{a_2,a_1} are just $\{a_1\}$ and $\{a_2\}$. Hence, under simple utility, an information structure yields excludability if and only if it has pairwise distinguishability.

More generally, though, the true value of excludability is that it permits a study of broad classes of preferences and information that are mutually sufficient for learning. We proceed to two such applications.

4. Applications

This section identifies informational conditions weaker than unbounded beliefs that guarantee adequate learning for two widely-used classes of preferences: with a one-dimensional state, preferences with single-crossing differences; and with a multi-dimensional state, (weighted) Euclidean preferences. These applications showcase how excludability can be fruitfully decoupled to obtain mutually sufficient classes of preferences and information.

4.1. Learning with Single-Crossing Preferences

In this subsection we take the state space to be totally ordered: for simplicity, $\Omega \subset \mathbb{R}$. A function $h : \Omega \to \mathbb{R}$ 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 $u : A \times \Omega \to \mathbb{R}$ have *single-crossing differences* (SCD) if for all a and a', the difference $u(a, \cdot) - u(a', \cdot)$ is single crossing.

SCD is an ordinal property closely related to a notion in Milgrom and Shannon (1994), but, following Kartik, Lee, and Rappoport (2022), the formulation is without an order on A. Ignoring indifferences, SCD requires that the preference over any pair of actions can only flip once as the state changes monotonically. If A is ordered, then SCD is implied by Milgrom and Shannon's (1994) notion or even a weaker one in Athey (2001).[18] SCD is thus

[18] Our treatment of indifferences is more permissive than Milgrom and Shannon (1994) and Kartik, Lee,

a property widely satisfied by economic models; in particular, it is assured by supermodularity of u.

The key informational condition will be 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 or assure that any state ω is distinguishable from all other states—not even any subset of states containing both some higher and some lower states than ω. Indeed, such a requirement would be incompatible with the widely-assumed MLRP, as was noted in Remark 1. What DUB does assure is that for any ω and prior μ with $\mu(\omega) > 0$, there are signals that make the posterior on the lower set $\{\omega' : \omega' < \omega\}$ arbitrarily small, and analogously for the upper set $\{\omega' : \omega' > \omega\}$. Equivalently, using Lemma 1, we can interpret DUB as saying 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 $\omega' > \omega$.[19]

A leading example of DUB information is *normal information*: signals are normally distributed on $\mathbb{R}$ with mean ω and fixed variance. For $\Omega \subseteq \mathbb{Z}$, one can see that DUB is satisfied here using Lemma 1: any state ω is distinguishable from its lower set because $\frac{f(s|\omega-1)}{f(s|\omega)} \to 0$ as $s \to \infty$ and for any $\omega' < \omega$, $\frac{f(s|\omega')}{f(s|\omega)} \leq \frac{f(s|\omega-1)}{f(s|\omega)}$; analogously, ω is distinguishable from its upper set by taking $s \to -\infty$.[20] Notably, however, no state can be simultaneously distinguished from both some lower one and some higher one. Figure 3a provides a graphical depiction. The solid blue curve therein is the set of posteriors obtained under normal information when $\Omega = \{1, 2, 3\}$ and the prior is μ. Posteriors are bounded away from certainty

and Rappoport (2022). A function $h : \{1, 2, 3\} \to \mathbb{R}$ given by $h(1) = 0$, $h(2) = 1$, and $h(3) = 0$ is single crossing per our definition but not according to the notions employed by Milgrom and Shannon (1994) and Kartik, Lee, and Rappoport (2022). SCD is, in fact, equivalent to saying that there exists some order on A with respect to which Athey's (2001) notion of "weak single-crossing property of incremental returns" is satisfied.

[19] In the context of auctions, Milgrom (1979) is concerned with distinguishing each state from its lower set (not also its upper set), and he notes in his Theorem 2 the relevant implication of our Lemma 1 for that kind of distinguishability. As he mentions, it is possible that any/all states are distinguishable from their lower sets but not their upper sets, or vice-versa, or they can be distinguishable from both or neither.

[20] This argument applies with only notational changes if $\Omega \not\subseteq \mathbb{Z}$, so long as every state has a maximum strictly lower state (if there are any strictly lower states) and a minimum strictly higher state (if there are any strictly higher states). Even absent this property (e.g., $\Omega = \mathbb{Q}$), normal information satisfies DUB by Lemma 2 in Subsection 4.2.

on state 2. Yet they can become arbitrarily certain about the relative probability of state 2 to 1 (as the curve approaches state 3) and, separately, about the relative probability of state 2 to 3 (as the curve approaches state 1): this is seen by projecting both the prior and the set of posteriors onto the relevant edges; for example, when the prior is μ' (the projection of μ onto the 2–3 edge) the set of posteriors is the entire 2–3 edge (the projection of the blue curve onto the 2–3 edge).

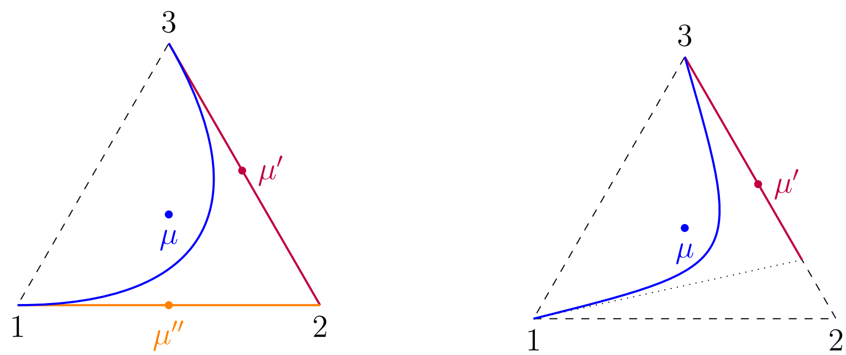

(a) DUB

(b) Not DUB

Figure 3: Illustration of DUB. For each belief, posteriors are shown in the corresponding color.

Beyond normal information, DUB is assured by any subexponential location-shift information structure, although we defer details to to Subsection 4.2. Alternatively, if the information structure has the MLRP, then DUB is equivalent to pairwise distinguishability.[21]

A failure of DUB was already seen in Figure 2: there, state 3 is not distinguishable from its lower set $\{1, 2\}$, nor is state 1 distinguishable from its upper set $\{2, 3\}$. Figure 3b displays a more nuanced failure. Here, states 1 and 3 are distinguishable from their complements, as the posteriors under the full-support prior μ, given by the solid blue curve, approach the relevant vertices. However, state 2 is not distinguishable from state 3: the posteriors for the prior μ' (the projection of μ onto the 2–3 edge), given by the solid red line (the projection of the blue curve onto the 2–3 edge), do not approach state 2's vertex. The reason is that unlike the case of normal information in Figure 3a, the blue curve in Figure 3b is not tangent to the 1–2 edge at the vertex 1.

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

We can now state:

Proposition 1. *If preferences have SCD and the information structure has DUB, then there is adequate learning. Conversely, if the information structure violates DUB, 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. SCD implies non-reversal of strict preferences over any pair of actions: for any pair of actions, a and a', either $\inf\{\omega : u(a,\omega) > u(a',\omega)\} \geq \sup\{\omega : u(a',\omega) > u(a,\omega)\}$ or $\inf\{\omega : u(a,\omega) < u(a',\omega)\} \geq \sup\{\omega : u(a',\omega) < u(a,\omega)\}$. DUB guarantees that all upper/lower sets are distinguishable from each other; more precisely, DUB is equivalent to saying 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 the sufficiency part of Proposition 1 follows from Theorem 2.

The argument for necessity of DUB in Proposition 1 is as follows. Assume DUB fails: say, some state ω^* cannot be distinguished from its lower states (it is analogous if "lower" is replaced by "upper"). Define 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 $u(a_2,\omega) = 1$. Any other actions can be assumed to be dominated and ignored. The failure of DUB implies that if a prior μ_0 is supported on $\omega \leq \omega^*$ and $\mu_0(\omega^*) > 0$ is small enough, then a_2 is never chosen after any signal. Any such prior is stationary but has inadequate knowledge. By Theorem 1, there is inadequate learning.[22]

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 there can be inadequate learning under DUB information absent SCD preferences. The following example illustrates.

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 $\varepsilon > 0$.[23] So any action ω is uniquely optimal in state ω

[22] Put differently, this argument shows that when DUB fails, there are SCD preferences such that excludability fails; the failure of adequate learning then follows from Theorem 2, or more precisely its infinite-state counterpart in the appendix, Theorem 2′. Note that the choice-set variation required by Theorem 2 is assured because SCD preferences are rich enough to stipulate any action as strictly dominated.

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

but worse than the safe action a^* in every other state. Plainly, SCD is violated.

Consider normal information. There are full-support priors such that the posterior probability of any state is uniformly bounded away from 1 across all the signals and states.[24] For any such prior, for small enough $\varepsilon > 0$, the safe action a^* is optimal after every signal. In other words, any such prior is stationary but has inadequate knowledge, and so Theorem 1 implies that there is inadequate learning. □

Remark 2. If Ω is finite, then DUB information yields adequate learning if there are distinct optimal actions at all states, i.e., for any $\omega \neq \omega'$, $c(\omega) \cap c(\omega') = \emptyset$. This follows from Theorem 1, because the only stationary beliefs are degenerate: for any belief μ, DUB implies that the posterior can be arbitrarily certain on the extreme states of μ's support. But Example 2 cautions that pairing DUB with distinct optimal actions at all states is not a robust principle for social learning. Furthermore, for some DUB information structures such as normal information, excludability fails whenever SCD is violated; hence, Theorem 2 implies there is inadequate learning at some (binary) choice set. Indeed, because SCD and DUB are a minimal pair of sufficient conditions for excludability in the terminology of Athey (2002)—if one condition fails, then the other can be satisfied though the preferences-information pair jointly fails excludability—Theorem 2 implies that SCD and DUB are a minimal pair of sufficient conditions for adequate learning at all choice sets.

4.2. Learning with Euclidean Preferences

In this subsection, we turn to a multidimensional state space and take $A, \Omega \subset \mathbb{R}^d$ for some integer $d \geq 1$. We view any $x \in \mathbb{R}^d$ as a column vector and denote its transposition by x' and its standard Euclidean norm by $\|x\|$.

Definition 5. Preferences represented by $u : A \times \Omega \to \mathbb{R}$ are *weighted Euclidean* if $u(a, \omega) = -l((a-\omega)'W(a-\omega))$, for some $d \times d$ symmetric positive definite matrix W and strictly increasing loss function $l : \mathbb{R}_+ \to \mathbb{R}_+$.

Weighted Euclidean preferences imply that in any state ω, an agent's indifference curves over actions are (rotated) ellipses, with actions inside the ellipse more preferred to those

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

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

As the last expression is the sum of a strictly positive decreasing function of s and a strictly positive increasing function of s, it is bounded away from 0 in s. The bound is independent of n because normal information is a location-shift family of distributions. Therefore, the posterior likelihood ratio is uniformly bounded away from 0, and hence, the posterior $\mu_s(n)$ is uniformly bounded away from 1.

outside.[25] A special case is the standard Euclidean preferences, $u(a,\omega) = -\sum_{i=1}^{d}(a_i - \omega_i)^2$, which has speherical indifference curves.

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

Definition 6. A location-shift information structure, or its standard density g, is *subexponential* if there are real numbers $p > 1$ 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.[26] 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 $M > 0$.

We can now state:

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

The result stems from Theorem 2, as the combination of subexponential location-shift information and Euclidean preferences yield excludability. To understand why excludability holds, notice first that the preferred sets for Euclidean preferences are *half spaces*, i.e. $\forall a_1 \neq a_2$, $\exists h \in \mathbb{R}^d$ and $c \in \mathbb{R}$ such that $\Omega_{a_1,a_2} = \{\omega : \ h \cdot \omega > c\}$. Excludability then follows from the lemma below.

Lemma 2. *For a subexponential location-shift information structure, $\{\omega : h \cdot \omega \geq c\}$ 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 $d = 1$ and the one-dimensional Laplace or double exponential standard density, $g(s) = (1/2)\exp(-|s|)$. This density is not subexponential,

[25] To ensure our maintained assumption of bounded utility, either the function l in Definition 5 must be bounded or the state and action spaces must be bounded.

[26] It would be more precise to refer to the density as *strictly* subexponential, as we require $p > 1$ in Definition 6, but we omit that adjective for brevity.

and indeed the conclusion of Lemma 2 fails: no two states are distinguishable from each other because, for example, if $\omega' < \omega < s$, then $f(s|\omega')/f(s|\omega) = g(s-\omega')/g(s-\omega) = \exp(\omega' - \omega)$ is independent of s.

We can provide an intuition for Lemma 2 by considering a bivariate normal information structure. Take an arbitrary hyperplane h, as illustrated in Figure 4. 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 on the right of the hyperplane from all the states on the left. Under normal information, the iso-density sets given any state are ellipses centered at the state. Given the location-shift structure, it is intuitively most difficult to distinguish ω_1 from states that are on (or arbitrarily close to) h, e.g. $\{\omega_2, \omega_3\}$. While it is natural to select a sequence of signals that grows unboundedly away from h (up and to the right), not all such sequences will work. Figure 4 shows how to construct a sequence of signals that distinguishes ω_1 from *all* states to the left of h. For a sequence of $c_n \to 0$, select s_n so that the iso-density ellipse at c_n given state ω_1 is tangent with the direction of h at s_n. Roughly, this selection implies that, regardless of c_n, the "distance" between s_n and ω_1 is boundedly (proportional to ε, the distance between ω_1 and the hyperplane h) smaller than that for any ω to the left of h. Due to the normal distribution being subexponential, as $c_n \to 0$ the likelihood ratio $\frac{g(s_n-\omega)}{g(s_n-\omega_1)} \to 0$ uniformly across ω to the left of h.[27,28]

More generally, when preferences are weighted Euclidean, for any two actions there is a hyperplane delineating when one is preferred to the other, and subexponential location-shift information ensures that the set of states on either side of the hyperplane can be distinguished from each other.

[27] Consider Figure 4. Since the signal distribution is bivariate normal, $g(s-\omega) \propto \exp[-(s-\omega)'\Sigma(s-\omega)]$ for some 2-by-2 matrix Σ. Consider signals on an iso-density ellipse of level c_n given state ω_1, i.e., $(s-\omega_1)'\Sigma(s-\omega_1) = c_n$. Let k denote the unit vector normal to hyperplane h. At signal s_n, the gradient of the level set is parallel to k, i.e., $2\Sigma(s_n-\omega_1) = 2\|\Sigma(s_n-\omega_1)\|k$. Hence,

$$\begin{aligned}
\ln\frac{g(s_n|\omega)}{g(s_n|\omega_1)} =&(s_n-\omega_1)'\Sigma(s_n-\omega_1) - (s_n-\omega)'\Sigma(s_n-\omega)\\
=&(s_n-\omega_1)'\Sigma(s_n-\omega_1) - (s_n-\omega_1+\omega_1-\omega)'\Sigma(s_n-\omega_1+\omega_1-\omega)\\
=&-2(\omega_1-\omega)'\Sigma(s_n-\omega_1) - (\omega_1-\omega)'\Sigma(\omega_1-\omega)\\
=&-2(\omega_1-\omega)'\|\Sigma(s_n-\omega_1)\|k - (\omega_1-\omega)'\Sigma(\omega_1-\omega)\\
\leq&-2\|\Sigma(s_n-\omega_1)\|\varepsilon - (\omega_1-\omega)'\Sigma(\omega_1-\omega).
\end{aligned}$$

As $c_n \to 0$, $\|\Sigma(s_n-\omega_1)\| \to \infty$, and so the ratio $\frac{g(s_n|\omega)}{g(s_n|\omega_1)} \to 0$ uniformly across ω to the left of h.

[28] In the bivariate normal case any such $c_n \to 0$ and associated sequence of signals s_n will suffice to distinguish ω_1 from the complementary half space. That is, it is only important that the sequence of signals grow unboundedly along the direction of the dotted line in Figure 4. For general subexponential distributions, this is not the case. But the proof of Lemma 2 shows that there is some sequence s_n that delivers the requisite distinguishability.

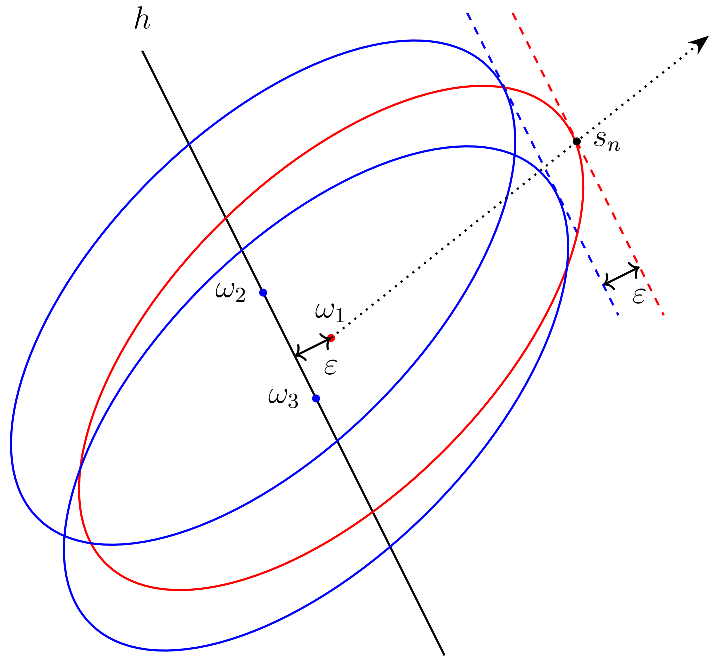

Figure 4: The logic underlying Lemma 2 for a bivariate normal standard density. We seek to distinguish ω_1 from the solid black line. The ellipses are level sets of the signal densities at the states ω_1, ω_2, and ω_3. As s_n grows along the dotted line, corresponding to lower iso-density levels, $\min\{f(s_n|\omega_2)/f(s_n|\omega_1), f(s_n|\omega_3)/f(s_n|\omega_1)\} \to 0$.

Remark 3. In the one-dimensional environment of Subsection 4.1, SCD is essentially equivalent to the half-space property of preferences just mentioned, and DUB is essentially equivalent to the distinguishability of half spaces from their complements. This perspective unifies both applications.

Remark 4. Although it is not precisely true that a subexponential standard density is necessary for a location-shift information structure to guarantee learning for all weighted Euclidean preferences, it can be shown that any superexponential standard density does preclude learning for some weighted Euclidean preferences.[29]

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 result that provides a welfare bound regardless of whether there is learning. Stating that result requires

[29] A density $g : \mathbb{R}^d \to \mathbb{R}$ is superexponential if there are $p \in (0, 1)$ and $M > 0$ such that $g(s) \geq \exp(-\|s\|^p)$ for all $\|s\| > M$.

some notation. Abusing notation, let

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

be the maximal utility an agent can get under belief μ. Writing μ_s for the posterior given a belief μ and signal s, let

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

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

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

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

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

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

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

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

We highlight that the only substantive requirement for this result is our maintained assumption that the observational network structure satisfies expanding observations. It is straightforward to see how Theorem 3 implies the "if" direction of Theorem 1. When all stationary beliefs have adequate knowledge, it holds for any distribution $\varphi \in \Phi^S$ that almost surely a correct action is taken. Thus, $u_*(\mu_0) = u^*(\mu_0)$, and we have adequate learning.

The argument behind Theorem 3 relies fundamentally on certain compactness and continuity. First, we establish in Lemma 4 of the appendix that Φ^{BP} is compact when the space of belief distributions is endowed with the Prohorov metric (which metrizes the weak topology). This is a consequence of Bayes-plausibility. Intuitively, for the case of countable states that our main text focuses on, although the prior μ_0 can be supported on an infinite set, it must concentrate an arbitrarily large mass on only finitely many states. More generally, for any $\delta > 0$, there exists a compact subset $\Omega' \subseteq \Omega$ such that $\mu_0(\Omega') \geq 1 - \delta$.[30] So any $\varphi \in \Phi^{BP}$ must put arbitrarily large probability on beliefs that put arbitrarily large probability on the compact set Ω' by Bayes plausibility, which yields compactness of Φ^{BP}. Next, owing to expected utility being continuous in beliefs, the improvement function $I(\varphi)$ is also continuous (Lemma 5 in the appendix), and thus uniformly continuous on Φ^{BP}. Consequently, $I(\varphi)$ achieves a minimum on any closed, hence compact, subset of Φ^{BP}.

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

We can then apply an *improvement principle*, as suggested by Banerjee and Fudenberg (2004) and developed by Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) and others. 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 interim belief distribution, φ. Her interim 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.[31] 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

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

[31] With 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.

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.[32] They consider a binary-state binary-action model. In that setting, they establish a welfare bound of "diffusion utility", which is the utility obtained by a hypothetical agent who observes an information structure that contains only the strongest signals (i.e., an "expert agent", in their terminology). Our cascade 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.[33] As Lobel and Sadler (2015) note, cascade utility and diffusion utility coincide in their binary-state binary-action model; but we note that in general, the former can be strictly higher.[34] Methodologically, Lobel and Sadler's (2015) 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. Our approach, by contrast, uses continuity of the improvement function and compactness of the set of Bayes-plausible belief distributions.

Beyond delivering conditions for learning, Theorem 3 is also useful when learning fails. For instance, we can use it to quantify how significantly failures of excludability reduce welfare. To make this point precise, say that for any $\varepsilon \in (0, 1/2)$ a set of states Ω' is ε-*distinguishable* from Ω'' if for any $\mu \in \Delta(\Omega' \cup \Omega'')$ with $\mu(\Omega') > \varepsilon$, there is a positive-measure set of signals S' such that $\mu(\Omega'|s) > 1-\varepsilon$ for all $s \in S'$. A utility function and an information structure jointly satisfy ε-*excludability* if Ω_{a_1,a_2} and Ω_{a_2,a_1} are ε-distinguishable from each other, for any pair of actions a_1, a_2. Note that ε-excludability implies ε'-excludability for all $\varepsilon' > \varepsilon$, and excludability is equivalent to ε-excludability for all $\varepsilon > 0$.

[32] In the complete network, where each agent observes every predecessor's action, Theorem 3 can be proved by a martingale argument, as done by Arieli and Mueller-Frank (2021, Lemma 1). They argue that the asymptotic "public belief" has to be stationary. However, this approach is inoperable in general networks because the public belief need not be a martingale and its convergence is not guaranteed.

[33] Lobel and Sadler's (2015) definition of diffusion utility is tailored to their binary-binary model. In general, we can define it as the highest utility an agent can obtain from any Bayes-plausible belief distribution that is supported on the set of feasible posteriors (i.e., those available in the given information structure and the prior); call the corresponding signal structure the expert signal structure. To see that this utility is lower than cascade utility, notice that this utility must be lower than from first drawing a posterior from an arbitrary Bayes-plausible stationary belief distribution and then drawing a signal from the expert signal structure (as this "combined signal" is Blackwell more informative than just the expert signal); but in the latter, the expert signal has no value by definition of stationary beliefs, and so the combined signal provides a utility equal to that from the (arbitrary) stationary belief distribution.

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

Proposition 3. *Let Ω be finite. For all $\varepsilon \in (0, 1/2)$, ε-excludability implies that in any equilibrium σ, $\liminf_n \mathbb{E}_{\sigma,\mu_0}[u_n] \geq u^*(\mu_0) - 2\overline{u}\frac{\varepsilon}{1-\varepsilon}|\Omega|$.*

The following example illustrates an application of Proposition 3.

Example 3. There are three states, $\omega \in \{1, 2, 3\}$, SCD preferences, and Laplace information:

$$f(s|\omega) = \frac{1}{2b}\exp\left(-\frac{|s-\omega|}{b}\right),$$

where $b > 0$ is a scale parameter; a smaller b corresponds to more precise information.

It is straightforward to verify that no two states can be distinguished from each other.[35] Therefore, not every stationary belief has adequate knowledge (so long as preferences are nontrivial), and by Theorem 1 there is inadequate learning.

Nonetheless, we claim that ε-excludability holds for any ε such that $\varepsilon > \frac{1}{1+\exp(\frac{1}{2b})}$. To see this, observe that since the information structure has MLRP and preferences satisfy SCD, we can focus on ε-distinguishing state 3 from 2 (or, equally, 2 from 1).[36] When $\varepsilon > \frac{1}{1+\exp(\frac{1}{2b})}$, we have $\frac{\varepsilon}{1-\varepsilon}\exp(1/b) > \frac{1-\varepsilon}{\varepsilon}$, so there exist signals that move the prior $(0, 1-\varepsilon, \varepsilon)$ to a posterior of at least $1-\varepsilon$ on state 3, which implies ε-distinguishability of state 3 from 2.

Proposition 3 implies that in any equilibrium, $\liminf \mathbb{E}_{\sigma,\mu_0}[u_n] \geq u^*(\mu_0) - 6\overline{u}\exp(-\frac{1}{2b})$. This quantitative welfare bound yields, in particular, convergence to the full-information utility $u^*(\mu_0)$ as $b \to 0$. □

6. Conclusion

This paper has studied a general model of sequential Bayesian social learning with observational networks. When there are multiple—i.e., more than two—states, learning depends jointly on society's preferences and information. While unbounded beliefs assures learning for all preferences, with multiple states it is incompatible with standard information structures, e.g., those with the monotone likelihood ratio property.

Fortunately, unbounded beliefs is unnecessary for learning in broad preference classes. We have identified a simple joint condition on preferences and information, *excludability*, that characterizes learning across all choice sets (Theorem 2). Excludability highlights how learning is driven by individual agents being able to rule out suboptimal actions rather

[35] For any pair of states $\omega \neq \omega'$, and any signal s, the likelihood ratio $f(s|\omega)/f(s|\omega') \leq \exp(2/b)$. Hence, no state can be distinguished from any other state.

[36] By MLRP, only arbitrarily large signals can distinguish a state from a lower state, and for large s the likelihood ratio $f(s|3)/f(s|2) < f(s|3)/f(s|1)$, so considering adjacent states is sufficient for ε-excludability.

than directly learn the correct action. In two applications, we have shown that excludability can be decoupled into separate conditions on preferences and information that are mutually sufficient for learning: one application concerned a one-dimensional environment with SCD preferences and DUB information (Proposition 1); and the other a multidimensional environment with weighted Euclidean preferences and subexponential location-shift information (Proposition 2).

Beyond excludability and its applications, we believe our Theorems 1 and 3, and the methodology they are based on, will be useful for other work on social learning. Theorem 1 provides a characterization of learning for a fixed choice set. Example 1 and Remark 2 illustrated how it can be used to deduce learning for a fixed choice set even when excludability fails. Theorem 3 is our key technical result, providing a welfare bound regardless of whether there is learning.

We close by commenting on two 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). A sufficient condition for our results in this case is that that expanding observations, i.e., condition (1), holds for the "induced network structure" $(\tilde{Q}_n)_{n\in\mathbb{N}}$ where each $\tilde{Q}_n$ is defined by first drawing a neighborhood B_n according to Q_n and then uniform-randomly drawing a single agent from B_n. See fn. 40 in the appendix for why this is sufficient. Interestingly, this 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 for all priors. To what extent this is true with multiple states is an interesting question for future research. We provide some incipient analysis in Appendix SA.1.

References

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

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

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

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

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

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

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

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

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

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

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

BLACKWELL, D. (1953): "Equivalent Comparisons of Experiments," *The Annals of Mathematical Statistics*, 24, 265–272.

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

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

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

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

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

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

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

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

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

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

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

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

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

KARTIK, N., S. LEE, AND D. RAPPOPORT (2022): "Single-Crossing Differences in Convex Environments," Unpublished.

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

LEHMANN, E. L. (1988): "Comparing Location Experiments," *The Annals of Statistics*, 16, 521–533.

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.

Persico, N. (2000): "Information Acquisition in Auctions," *Econometrica*, 68, 135–148.

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

Rosenberg, D. and N. Vieille (2019): "On the Efficiency of Social Learning," *Econometrica*, 87, 2141–2168.

Smith, L. and P. Sørensen (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68, 371–398.

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

A. Backbone Results

We prove our backbone results—Theorem 1, Theorem 2, and Theorem 3—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;[37,38]
- The utility function $u(a, \omega)$ has absolute value uniformly bounded by $\overline{u}$ 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 $\{1, 2, \ldots, n-1\}$, independent across n, independent of the state ω, and independent of any private signals.

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

With these elements, we can define an overarching probability space over all realizations of the state, signals, observation neighborhoods, and actions. An overarching probability measure is determined uniquely by the prior μ_0, the signal structure F, the network structure Q, and a strategy profile σ. We will refer to this overarching probability measure by $\mathbb{P}_{\sigma,\mu_0}$, or sometimes just $\mathbb{P}$ for short. The agents' beliefs are defined as the regular conditional probabilities of this overarching probability measure, which are guaranteed to exist and are unique almost everywhere. Appendix SA.2 elaborates.

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

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

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

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

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

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

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 $\varphi \in \Phi^{BP}$ has to put a large probability on a compact subset of $\Delta\Omega$.

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

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

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

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

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

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

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

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

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

Lemma 4. *Φ^{BP} is compact.*

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

Take any open set $W \in \mathcal{B}(\Omega)$. For any $\mu_k \to \mu$, 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 $\mathcal{B}(\Omega)$.

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

A.2. Continuity of Various Functions

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Now consider the function u. Its continuity follows from

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

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

Lastly, suppose we establish that $\Gamma^F := \left\{ \int_S u(\beta(s),\omega)\,\mathrm{d}F(s|\omega) \right\}_{\beta:S\to A}$, as a family of functions of ω indexed by β, is pointwise equicontinuous.[39] Then, as Γ^F is uniformly bounded, Claim 1 implies that $u^F(\mu)$ is continuous, proving the lemma.

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

$$\begin{aligned} &\left| \int_S u(\beta(s),\omega)\,\mathrm{d}F(s|\omega) - \int_S u(\beta(s),\omega')\,\mathrm{d}F(s|\omega') \right| \qquad (3) \\ \le &\left| \int_S (u(\beta(s),\omega) - u(\beta(s),\omega'))\,\mathrm{d}F(s|\omega) \right| + \left| \int_S u(\beta(s),\omega')\,\mathrm{d}F(s|\omega) - \int_S u(\beta(s),\omega')\,\mathrm{d}F(s|\omega') \right|. \end{aligned}$$

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

Fix any ω and any $\varepsilon > 0$. Since $\{u(a,\omega)\}_{a\in A}$ is pointwise equicontinuous, there exists δ_1 such that $d(\omega',\omega) < \delta_1$ implies the first term on the right-hand side of inequality (3) to be smaller than $\varepsilon/2$ (regardless of $\beta(s)$). The second term is smaller than $2\overline{u}d_{TV}(F(\cdot|\omega), F(\cdot|\omega'))$ (where TV represents total variation), and by the continuity assumption of the information structure, there exists $\delta_2 > 0$ such that $d(\omega',\omega) < \delta_2$ implies $d_{TV}(F(\cdot|\omega), F(\cdot|\omega')) < \varepsilon/4\overline{u}$. Therefore, if $d(\omega',\omega) < \min\{\delta_1, \delta_2\}$, then the right-hand side of inequality (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, define $u(\varphi) := \mathbb{E}_\varphi[u(\mu)]$, $I(\varphi) := \mathbb{E}_\varphi[I(\mu)]$, and $G(\varphi) := \mathbb{E}_\varphi[G(\mu)]$ as the corresponding functions over distributions of beliefs. Since $u(\mu)$, $I(\mu)$, and $G(\mu)$ are continuous, so are $u(\varphi)$, $I(\varphi)$, and $G(\varphi)$.

A.3. Proofs for Backbone Results

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

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

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

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

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

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

We claim that, for any agent $n \geq N_k$, $\mathbb{E}_{\sigma,\mu_0}[u_n] \geq \alpha_k := \min\{u_* - \gamma(\varepsilon), \frac{kI(\underline{\varphi})}{2} - \overline{u}\}$. Since $\alpha_0 = -\overline{u}$, clearly $\mathbb{E}_{\sigma,\mu_0}[u_n] \geq \alpha_0$ for any $n \geq N_0$. Suppose the claim holds for all agents $n' \geq N_{k-1}$. Take any agent $n \geq N_k$. Agent n's neighborhood is drawn independently of everything that has happened before, so conditional on agent n observing an agent $n' \geq N_{k-1}$, even without her private signal agent n can achieve a utility of at least α_{k-1} by imitating agent n'. Hence, $u(\varphi_n) \geq (1-\delta) \cdot \alpha_{k-1} + \delta \cdot (-\overline{u})$.[40] If $\varphi_n \in (\Phi^S)^\varepsilon$, then by definition $u(\varphi_n) \geq u_* - \gamma(\varepsilon)$, and thus $\mathbb{E}_{\sigma,\mu_0}[u_n] \geq u(\varphi_n) \geq u_* - \gamma(\varepsilon) \geq \alpha_k$. If $\varphi_n \notin (\Phi^S)^\varepsilon$, then agent n can improve her utility by at least $I(\underline{\varphi})$, and so

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

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

[40] If 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.

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

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

which implies $u(\varphi) = u^*(\mu_0)$. As a result, $u_*(\mu_0) = \inf_{\varphi \in \Phi^S} u(\varphi) = u^*(\mu_0)$. It follows from Theorem 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 $\mathbb{E}_{\sigma,\mu_0}[u_n] \to u^*(\mu_0)$. 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 $\Omega^n_{a_1,a_2} := \{\omega : u(a_1, \omega) - u(a_2, \omega) > \frac{1}{n}\}$.

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 $\Omega^n_{a_2,a_1}$ for some n.*

Note that when Ω is finite, or the utility difference between between any pair of actions is bounded away from zero, a failure of excludability is equivalent to 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′. (<u>First statement</u>) 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^n_{a',a^*} are open sets for any a' and n. Since Ω is Polish, it is second-countable and hence has a countable basis. Therefore, each open set Ω^n_{a',a^*}, and hence the open set $\cup_{a'} \Omega_{a',a^*} (= \cup_{a'} \cup_n \Omega^n_{a',a^*})$, 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^n_{a',a^*} for some a' and n has strictly positive measure, i.e., $\mu(\Omega^n_{a',a^*}) > 0$.

Now denote $\mu'(\cdot) := \mu(\cdot|\Omega_{a^*,a'} \cup \Omega^n_{a',a^*})$ as the corresponding conditional probability. Since Ω_{a',a^*} is distinguishable from $\Omega_{a^*,a'}$ by excludability, so is Ω^n_{a',a^*}.[41] Therefore, for any $\varepsilon > 0$ there exists a set of signals S' such that $\mathbb{P}_{\mu'}(S') > 0$ and $\mu'_s(\Omega^n_{a',a^*}) > 1-\varepsilon$ for all $s \in S'$. The utility improvement upon observing any $s \in S'$ by switching from a^* to a' is therefore bounded below by $(\frac{1}{n}(1-\varepsilon) - 2\overline{u}\varepsilon)\mu_s(\Omega_{a^*,a'} \cup \Omega^n_{a',a^*})$, as the expected improvement on $\Omega\backslash(\Omega_{a^*,a'}\cup\Omega^n_{a',a^*})$ is nonnegative. For small $\varepsilon > 0$, $(\frac{1}{n}(1-\varepsilon)-2\overline{u}\varepsilon) > 0$, and hence, integrating over S', the ex-ante improvement is bounded below by $(\frac{1}{n}(1-\varepsilon) - 2\overline{u}\varepsilon)\mathbb{P}_{\mu'}(S')\mu(\Omega_{a^*,a'} \cup \Omega^n_{a',a^*}) > 0$. It follows that $I(\mu) > 0$, and thus μ is not stationary.

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

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

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

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

B. Applications

To the general setup of Appendix A, we now add the assumption that $\Omega \subset \mathbb{R}$ is countable, endowed with the discrete metric, and $F(\cdot|\omega)$ are absolutely continuous with respect to each to other, and so there are densities $f(\cdot|\omega) > 0$.

[41] In fact, excludability is equivalent to: $\Omega^n_{a_1,a_2}$ 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. Half-Space Preferences & Location-Shift Information

In this section we assume $A \subseteq \mathbb{R}^d$ and a countable $\Omega \subset \mathbb{R}^d$. The proof of Lemma 2 below is more involved than the intuition given in the main text using Figure 4, 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.

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 $\frac{g(s')}{g(s)} \geq 2$. From 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 $\lim_{i\to\infty} \|s_i\|_h = \infty$. It holds that for all s_i,

$\frac{g(s_i)}{g(\overline{s})} \geq \frac{g(s_1)}{g(\overline{s})} \varepsilon^{(k_{s_{i-1}}+\cdots+k_{s_1})\|\overline{s}\|_h} = \frac{g(s_1)}{g(\overline{s})} \varepsilon^{(\|s_i\|_h - \|\overline{s}\|_h - 1)\|\overline{s}\|_h}$, which in turn implies that

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

However, since g is subexponential, and $\|s_i\|_h$ and $\|s_i\|$ are of the same order when $\|s_i\|_h$ is large, there is $p > 1$ such that for all large enough i,

$$\log(g(s_i)) < -(\|s_i\|_h)^p. \tag{5}$$

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

Proof of Lemma 2. We only prove that $\{\omega : h\cdot\omega \geq c\}$ is distinguishable from $\{\omega : h\cdot\omega < c\}$. The other direction is similar because any state in $\{\omega : h \cdot \omega < c\}$ will have a neighborhood that is strictly separated from $\{\omega : h \cdot \omega \geq c\}$.

Let $\|x\|_h := h\cdot x - c$ be the "signed distance" between x and the hyperplane $\{z : h\cdot z = c\}$. 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 $\lim_{i\to\infty} \|s_i^*\|_h = \infty$.

As noted after Definition 1, it is sufficient to prove that any $\overline{\omega} \in \{\omega : h \cdot \omega \geq c\}$ is distinguishable from $\{\omega : h \cdot \omega < c\}$.[42] So take any such $\overline{\omega}$ and μ with $\mu(\overline{\omega}) > 0$. Define $\overline{s}_i := s_i^* + \overline{\omega}$. It follows that for all i,

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

and

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

[42] We note that this uses the assumption of countable states.

and thus,

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

The last expression can be taken arbitrarily small because $\|s^*_{i-1}\|_h \to \infty$ as $i \to \infty$. Finally, by uniform continuity of g there exists a neighborhood of signals $\overline{S}_i$ around $\overline{s}_i$ over which the last inequality above continues to hold. Therefore, $\overline{\omega}$ is distinguishable from $\{\omega : h\cdot\omega < c\}$. *Q.E.D.*

B.3. Other Proofs

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 $1 \in \operatorname{Supp}\mu(\omega'|\cdot)$.

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.*

Proof of Proposition 3. Take any stationary belief μ, and let a be an optimal action at belief μ. For each state ω, take any $a_\omega \in c(\omega)$, and consider $\mu_\omega(\cdot) := \mu(\cdot|\{\omega\}\cup\Omega_{a,a_\omega})$. If $\mu_\omega(\omega) \le \varepsilon$, then $\mu(\omega) \le \varepsilon$, so $(u(a_\omega,\omega) - u(a,\omega))\mu(\omega) \le 2\overline{u}\varepsilon$.

Consider the other case of $\mu_\omega(\omega) > \varepsilon$. For any $s \in S$, because $u(a,\omega') - u(a_\omega,\omega') \le 0$ for each $\omega' \notin \Omega_{a,a_\omega}$, and μ is stationary,

$$\sum_{\omega'\in\{\omega\}\cup\Omega_{a,a_\omega}} (u(a,\omega') - u(a_\omega,\omega'))\mu(\omega'|s) \ge \sum_{\omega'\in\Omega}(u(a,\omega') - u(a_\omega,\omega'))\mu(\omega'|s) \ge 0.$$

Then,

$$\begin{aligned}(u(a_\omega,\omega)-u(a,\omega))\mu_\omega(\omega|s) &\leq \sum_{\omega'\in\Omega_{a,a_\omega}}(u(a,\omega')-u(a_\omega,\omega'))\mu_\omega(\omega'|s)\\ &\leq 2\overline{u}\left(\sum_{\omega'\in\Omega_{a,a_\omega}}\mu_\omega(\omega'|s)\right)=2\overline{u}(1-\mu_\omega(\omega|s)).\end{aligned}$$

By ε-excludability, there exists a positive-measure set of signals S' such that, for any $s\in S'$, $\mu_\omega(\omega|s)>1-\varepsilon$, which implies that $u(a_\omega,\omega)-u(a,\omega)\leq 2\overline{u}\frac{\varepsilon}{1-\varepsilon}$.

In either case ($\mu_\omega(\omega)\leq\varepsilon$ or $\mu_\omega(\omega)>\varepsilon$), we have $(u(a_\omega,\omega)-u(a,\omega))\mu(\omega)\leq 2\overline{u}\frac{\varepsilon}{1-\varepsilon}$. Since Ω is finite,

$$\sum_{\omega\in\Omega}(u(a_\omega,\omega)-u(a,\omega))\mu(\omega)\leq 2\overline{u}\frac{\varepsilon}{1-\varepsilon}|\Omega|.$$

Namely, the utility gap $u^*(\mu)-u(\mu)\leq 2\overline{u}\frac{\varepsilon}{1-\varepsilon}|\Omega|$, for any stationary belief μ.

Finally, for any $\varphi\in\Phi^S$,

$$u^*(\mu_0)-u(\varphi)=\mathbb{E}_\varphi[u^*(\mu)-u(\mu)]\leq 2\overline{u}\frac{\varepsilon}{1-\varepsilon}|\Omega|.$$

By taking infimum of $u(\varphi)$ across $\varphi\in\Phi^S$, we obtain $u_*(\mu_0)\geq u^*(\mu_0)-2\overline{u}\frac{\varepsilon}{1-\varepsilon}|\Omega|$, and subsequently by invoking Theorem 3, we conclude that in any equilibrium σ, $\liminf_n \mathbb{E}_{\sigma,\mu_0}[u_n]\geq u^*(\mu_0)-2\overline{u}\frac{\varepsilon}{1-\varepsilon}|\Omega|$. *Q.E.D.*

Supplementary Appendices (Not For Publication)

SA.1. Learning at a Fixed Prior

For tractability, our discussion in this appendix assumes a complete network.

The issue. The definition of adequate learning we adopted in Section 2 requires that there is learning for all priors. In general, one may be interested in whether there is adequate learning at some given (full-support) prior.[43] Of course, our sufficient conditions—e.g., Theorem 2's excludability—remain sufficient, but fixing a prior raises the question of whether the conditions are necessary. With only two states, the distinction between some prior and all priors is immaterial: if adequate learning fails at any prior, then the only adequate-knowledge beliefs are those with certainty on some state, and there is an open ball of stationary beliefs around certainty on one of the states; hence, given any full-support prior, there cannot be a belief path that converges to certainty on that state, implying a failure of adequate learning at all full-support priors.

However, with multiple states, a failure of adequate learning at some prior does not imply an open ball of stationary beliefs around any adequate-knowledge belief. To illustrate, consider Figure SA.1. Action a^* is optimal at states 2 and 3 while $\underline{a}$ is optimal at state 1. Adequate learning fails when the prior is μ because μ, which has support $\{1,2\}$, is stationary but has inadequate knowledge.[44] Yet there is no open ball of stationary beliefs: no full-support belief is stationary because the optimal actions are distinct at the extreme states 1 and 3, and the extreme states are distinguishable from their complements. This raises the possibility that there is learning at some—or even all—full-support priors, with on-path sequences of beliefs (which necessarily have full support at every finite time) converging almost surely to adequate-knowledge beliefs without ever hitting any stationary belief (all of which have non-full-support).

Some partial analysis. While we are unable to characterize learning at a fixed prior in general, we provide some partial analysis below that we hope will be useful for future research. We focus on obtaining an analog of Theorem 2 for a fixed prior.

[43] To be complete: there is *adequate learning at prior* μ_0 if for every equilibrium strategy profile σ, $\mathbb{E}_{\sigma,\mu_0}[u_n] \to u^*(\mu_0)$. Adequate (or inadequate) learning at a prior for a choice set is defined analogously.

[44] In this example, preferences have SCD. But, consistent with Proposition 1, DUB is violated because state 2 is not distinguishable from state 1 or state 3.

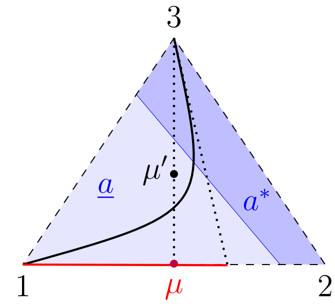

Figure SA.1: Preference regions among the actions $\underline{a}$ and a^* shaded in blues. Under belief μ', posteriors are given by the black curve, while under belief μ, posteriors are given by the red line.

First, we provide a lemma (Lemma SA.1) that connects the existence of certain on-path histories to distinguishability. Second, we conjecture a result (Conjecture SA.1) and show in Proposition SA.1 that, if the result is true, it combines with the lemma to deliver a fixed-prior analog of Theorem 2. Third, we show that the conjecture is true in a class of problems (Claim SA.1).

Lemma SA.1. *Take an aribtrary $\omega^* \in \Omega$ and set of states $\Omega' \subseteq \Omega \backslash \{\omega^*\}$. State ω^* is distinguishable from Ω' if there exist an equilibrium under a full-support prior and a history of actions h^∞ such that $\mathbb{P}(h^\infty|\omega^*) = 0$, and $\mathbb{P}(h^\infty|\omega)$ is bounded away from 0 across $\omega \in \Omega'$.*

The lemma ties the asymptotic probabilities of on-path histories to the information structure of an individual agent. The formal proof of the lemma is provided at the end of this appendix, but to see the intuition, suppose that the relevant h^∞ in the hypothesis of the lemma is some eventual herd on some action $a \in A$, i.e., $h^\infty = \{h^m, a, a, a, , ...\}$ with h^m some finite subhistory. If h^∞ has 0 probability in state ω^*, it must be because in an infinite number of periods, agents have positive probability of obtaining signals which overturn the herd on a, i.e., result in them taking some other action than a. However, the probability of this history is positive for states in Ω'. This means that the probability of signals that overturn the herd must vanish over time at a fast enough rate in ω^*, but either not vanish or vanish at a slow enough rate in each $\omega \in \Omega'$. In particular, there must exist overturning signals whose probability gets arbitrarily large in state ω^* relative to those in every $\omega \in \Omega'$, which means ω^* is distinguishable from Ω'.

Conjecture SA.1. *Take any $a_1, a_2 \in A$, any full-support prior, and any equilibrium. If there is adequate learning at that prior, choice set $\{a_1, a_2\}$, and equilibrium,*[45] *then*

$$\exists h^\infty \text{ and } \varepsilon > 0: \ \mathbb{P}(h^\infty|\omega) > \varepsilon, \forall \omega \in \Omega_{a_1,a_2}. \tag{SA.1}$$

[45] That is, under the given prior μ_0 and equilibrium σ, $\mathbb{E}_{\sigma,\mu_0}[u_n] \to u^*(\mu_0)$.

The conjecture says that given any full-support prior, any binary choice set $\{a_1, a_2\}$, and any equilibrium in which there is adequate learning, we can find a single history that occurs with probability bounded away from 0 in all states in which a_1 is strictly preferred (and analogously, a different history for the states in which a_2 is strictly preferred). To appreciate the conjecture, let us focus for discussion on the case of finite states, nontrivial information, and nontrivial preferences. First note that if Ω_{a_1,a_2} is a singleton—as is the case with binary states—then it is straightforward that there is such a history, as there is a herd almost surely and every herd begins at some finite time. When Ω_{a_1,a_2} is not a singleton, given adequate learning, the same logic shows that for each state in Ω_{a_1,a_2}, there is a history that has positive probability in that state, namely one with a herd on a_1. But Conjecture SA.1 demands more: a single history that has positive probability in all states in Ω_{a_1,a_2}. Nonetheless, the conjecture seems intuitive: (up to tie-breaking issues) it would be surprising for every infinite history that has positive probability in some $\omega \in \Omega_{a_1,a_2}$ to have zero probability in some other $\omega' \in \Omega_{a_1,a_2}$, given that agents' have the same ordinal preferences over the binary actions in both ω and ω'. For instance, consider a fully-informative information structure and any nontrivial preferences. Clearly, given any choice set $\{a_1, a_2\}$, there are only two possible histories: either a_1 in every period or a_2 in every period. The former has probability 1 in each $\omega \in \Omega_{a_1,a_2}$, and the latter has probability 1 in each $\omega \in \Omega_{a_2,a_1}$ and so Conjecture SA.1 holds. Even though individuals' private information distinguishes states perfectly, the public history does not.

Proposition SA.1. *If Conjecture SA.1 is true, then not only does excludability imply adequate learning at every prior for every choice set, but moreover, if excludability fails, then there exists a choice set with inadequate learning at every full-support prior.*

Proof. That excludability implies adequate learning at every prior for every choice set is implied by Theorem 2, with no need to invoke Conjecture SA.1. So we only prove the second portion of the proposition, doing so by contraposition.

To that end, assume that for every choice set, there is some full-support prior at which there is adequate learning in some equilibrium. For every binary choice set $\{a_1, a_2\}$, Conjecture SA.1 implies the existence of a history h^∞ satisfying (SA.1) at the full-support prior and equilibrium at which there is adequate learning. Since there is adequate learning, eventually all agents must be taking a_1 in h^∞, which implies that $\mathbb{P}(h^\infty|\omega^*) = 0$ for each $\omega^* \in \Omega_{a_2,a_1}$. Then, taking $\Omega' = \Omega_{a_1,a_2}$ in Lemma SA.1 yields that ω^* is distinguishable from Ω_{a_1,a_2}. Since a_1, a_2 and $\omega^* \in \Omega_{a_1,a_2}$ are arbitrary, there is excludability. *Q.E.D.*

We have not been able to establish Conjecture SA.1 in general. However, we are able to establish it when preferences satisfy SCD and the information structure satisfies the strict MLRP (assuming, only for convenience, that the state space is discrete):

Claim SA.1. *Assume $\Omega \subseteq \mathbb{Z}$. If preferences satisfy SCD and the information structure satisfies the strict MLRP, then Conjecture SA.1 is true.*

The proof is at the end of this appendix. Combining Claim SA.1 and Proposition SA.1, we see that under a complete network, the signal structure and preferences in Figure SA.1 entail inadequate learning at every full-support prior, such as μ' in the figure. Note that the figure's signal structure satisfies strict MLRP because the black curve in Figure SA.1 is concave vis-à-vis the 1–3 edge and approaches the 1 and 3 vertices.

Omitted Proofs

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

$$\int_{\Omega'} f(s|\omega)\,\mathrm{d}\tilde{\mu}(\omega) \geq z f(s|\omega^*). \tag{SA.2}$$

Suppose there exist an equilibrium σ under a full support prior and history h^∞ such that $\mathbb{P}(h^\infty|\omega^*) = 0$, and $\mathbb{P}(h^\infty|\omega)$ is bounded away from 0 across $\omega \in \Omega'$. Let a^n be the action taken by agent n along h^∞ and $A^{-n} := A\backslash\{a^n\}$. Let $\mathbb{P}(a^n|h^n,\omega) := \int_S \sigma(a^n|s,h^n) f(s|\omega)\,\mathrm{d}s$ be the probability that agent n plays action a^n when the state is ω and the sub-history is h^n. It holds that:

$$\begin{aligned}
\sum_{n=1}^{\infty} \log(1 - z\mathbb{P}(A^{-n}|h^n,\omega^*)) &\geq \sum_{n=1}^{\infty} \log\left(1 - \int_{\Omega'} \mathbb{P}(A^{-n}|h^n,\omega)\,\mathrm{d}\tilde{\mu}(\omega)\right) && \text{(using (SA.2))}\\
&= \sum_{n=1}^{\infty} \log\left(\int_{\Omega'} \mathbb{P}(a^n|h^n,\omega)\,\mathrm{d}\tilde{\mu}(\omega)\right)\\
&\geq \sum_{n=1}^{\infty} \int_{\Omega'} \log(\mathbb{P}(a^n|h^n,\omega))\,\mathrm{d}\tilde{\mu}(\omega) && \text{(by Jensen's inequality)}\\
&= \int_{\Omega'} \sum_{n=1}^{\infty} \log(\mathbb{P}(a^n|h^n,\omega))\,\mathrm{d}\tilde{\mu}(\omega) && \text{(by Tonelli's theorem)}\\
&= \int_{\Omega'} \log\left(\prod_{n=1}^{\infty} \mathbb{P}(a^n|h^n,\omega)\right)\mathrm{d}\tilde{\mu}(\omega)\\
&> -\infty \qquad \text{(as } \log \mathbb{P}(h^\infty|\omega) \text{ is bounded across } \omega \in \Omega'). && \text{(SA.3)}
\end{aligned}$$

Below we will invoke the fact that for arbitrary sequences (S_n) and (T_n) and constant $c > 0$, if $\lim_{n\to\infty} \frac{S_n}{T_n} = c > 0$ and $\sum_n S_n < \infty$, then $\sum_n T_n < \infty$.[46] Let $S_n = -\log(1 - z\mathbb{P}(A^{-n}|h^n, \omega^*))$ and $T_n = -\log(1-\mathbb{P}(A^{-n}|h^n, \omega^*))$. Note that $\lim_{n\to\infty} \frac{S_n}{T_n} = z \in (0, 1)$ because $\lim_{x\to 0} \frac{\log(1-zx)}{\log(1-x)} = z$ and (SA.3) implies $\lim_{n\to\infty} \mathbb{P}(A^{-n}|h^n, \omega^*) = 0$. The aforementioned mathematical fact implies that

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

As $\mathbb{P}(a^n|h^n, \omega^*) = 1 - \mathbb{P}(A^{-n}|h^n, \omega^*)$, it further follows that $\prod_{n=1}^{\infty} \mathbb{P}(a^n|h^n, \omega^*) > 0$, which contradicts $\mathbb{P}(h^\infty|\omega^*) = 0$. *Q.E.D.*

Proof of Claim SA.1. Take any information structure f with strict MLRP, and a utility function u that has SCD. Then, take an equilibrium σ under a full support prior μ and a binary choice set $\{a_1, a_2\}$. Since u has SCD, Ω_{a_1,a_2} and Ω_{a_2,a_1} are either an upper and lower set or the other way around. We consider the case that a_2 is preferred in higher states and a_1 is preferred in lower states. We omit an analogous proof for the other case.

We observe that $\mathbb{P}(a_1|h^n, \omega) := \int_S \sigma(a_1|h^n, s) f(s|\omega)\, \mathrm{d}s$, the probability that agent n plays action a_1 given any finite history h^n, decreases in ω. First, the probability $\sigma(a_1|h^n, s)$ decreases in s. The interim belief $\mu(\cdot|h^n)$ has full support, so strict MLRP of the information structure implies that $\forall s < s'$, the posterior $\mu(\cdot|h^n, s')$ strictly monotone likelihood-ratio dominates $\mu(\cdot|h^n, s)$. By Theorem 2 of Athey (2002),

$$D(s) := \sum_{\omega} (u(a_2, \omega) - u(a_1, \omega))\, \mu(\omega|h^n, s)$$

is strictly single crossing in s, i.e., $D(s) \geq 0 \implies D(s') > 0, \forall s' > s$. Hence, $\sigma(a_1|h^n, s)$ is decreasing in s. Since f satisfies strict MLRP, $\mathbb{P}(a_1|h^n, \omega)$ is decreasing in ω.

Next, suppose there is adequate learning. So for each state $\omega \in \Omega_{a_1,a_2}$, there is an infinite history with a herd on a_1, $h^\infty = (\ldots, a_1, a_1, \ldots)$ that occurs with positive probability in ω. In particular, if we let $\tilde{\omega} = \max \Omega_{a_1,a_2}$, then for any finite sub-history h^n of h^∞,

$$\forall \omega \leq \tilde{\omega} : \mathbb{P}(a_1|h^n, \omega) \geq \mathbb{P}(a_1|h^n, \tilde{\omega}) > 0,$$

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

and since h^∞ has positive probability at $\tilde{\omega} \in \Omega_{a_1,a_2}$, it follows that

$$\forall \omega \leq \tilde{\omega} : \prod_{n=1}^{\infty} \mathbb{P}(a_1 | h^n, \omega) \geq \prod_{n=1}^{\infty} \mathbb{P}(a_1 | h^n, \tilde{\omega}) > 0. \tag{SA.4}$$

This means $\mathbb{P}(h^\infty|\omega)$ is uniformly bounded away from 0 for $\{\omega : \omega \leq \tilde{\omega}\}$. Since $\Omega_{a_1,a_2} \subseteq \{\omega : \omega \leq \tilde{\omega}\}$, this establishes Conjecture SA.1. *Q.E.D.*

SA.2. Overarching Probability Space and Beliefs

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

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

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

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

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

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