Observational Learning with Ordered States*

Navin Kartik[†] SangMok Lee[‡] Daniel Rappoport[§]

April 30, 2021

Abstract

When does society eventually learn the truth, or underlying state, via sequential observational learning? This paper develops the interplay of preferences satisfying single-crossing differences (SCD) and a new informational condition, directionally unbounded beliefs (DUB). SCD preferences and DUB information are a jointly minimal pair of sufficient conditions for learning. When there are more than two states, DUB is weaker than unbounded beliefs, which characterizes learning for all preferences (Smith and Sørensen, 2000). Unlike unbounded beliefs, DUB is compatible with the monotone likelihood ratio property, and satisfied, for example, by normal information.

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

*We thank Nageeb Ali, Marina Halac, Elliot Lipnowski, José Montiel Olea, Xiaosheng Mu, Evan Sadler, Lones Smith, and Peter Sørensen for useful comments. 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@gmail.com

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

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

1. Introduction

This paper concerns the classic sequential observational 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 homogenous preferences over her own action and the state (e.g., all prefer products of higher quality). But agents act in sequence, each having access only to her own private information about the state and the observational history of predecessors' actions. The central economic question is about asymptotic learning: does society eventually learn the true state from the history of individual actions (e.g., will the highest quality product eventually prevail)?

In their fundamental work, Smith and Sørensen (2000) identify *unbounded signals/beliefs* as sufficient to guarantee (asymptotic) learning for all preferences.[1] To illustrate, consider ex-ante equiprobable states, $\omega \in \Omega = \{1, 2\}$, a finite set of actions A with $\Omega \subseteq A \subset \mathbb{R}$, and each agent's utility from taking action a in state ω given by $u(a, \omega) = -(a - \omega)^2$. Private signals drawn independently from normal distributions with mean ω and fixed variance satisfy unbounded beliefs. The key observation is that, no matter the history of actions, there will always be an agent who gets a sufficiently strong signal in favor of the true state (a sufficiently low signal if $\omega = 1$ and a sufficiently high one if $\omega = 2$) that leads her to take the optimal action $a = \omega$. Thus, society can never get stuck at a belief at which agents are taking incorrect actions. This observation, combined with convergence arguments, is enough to imply here that there is learning and eventually all agents take the optimal action $a = \omega$.

Many, if not most, applications of observational learning (e.g., product choice) ought to incorporate more than two states. Consider amending the above example to have states $\Omega = \{1, 2, 3\}$. The normal information structure no longer satisfies unbounded beliefs: there are no signals that can make one arbitrarily sure about state 2. Yet Smith and Sørensen's (2000) sufficient condition to rule out learning, *bounded signals/beliefs*, is not met either: arbitrarily low signals make one arbitrarily sure about state 1, while arbitrarily high signals make one arbitrarily sure about state 3. What, then, can one deduce about learning in this canonical example?

Our main result, Theorem 1 in Section 4, implies that there is learning. The reason is that the normal information structure satisfies a property we dub *directionally unbounded beliefs* (DUB), and the quadratic-loss preferences satisfy *single-crossing differences* (SCD). SCD is a widely-invoked property on preferences in economics; its implications for comparative

[1] Smith and Sørensen's (2000) model allows for heterogenous preference types; throughout this paper, we are concerned with a single preference type—what they refer to as "the traditional case".

statics were developed by Milgrom and Shannon (1994). We are not aware, however, of it having been applied to learning. Following Kartik, Lee, and Rappoport (2019), the version of SCD we use requires that agents' preferences over any pair of actions do not reverse more than once as the state changes monotonically. SCD is implied by supermodularity of the utility function, which holds for quadratic loss. DUB, on the other hand, appears to be a new condition on information structures. Like SCD, DUB is formulated for an ordered state space. Roughly speaking, it requires that for any state ω there exist both: (i) signals that make one arbitrarily sure about ω given any prior with support the lower set $\{\omega' : \omega' \leq \omega\}$; and (ii) signals that make one arbitrarily sure about ω given any prior with support the upper set $\{\omega' : \omega' \geq \omega\}$. For the normal information structure, requirements (i) and (ii) are met for any state by arbitrarily high and arbitrarily low signals, respectively.

Theorem 1 establishes that DUB information and SCD preferences are jointly sufficient for learning, or more precisely, adequate learning à la Aghion, Bolton, Harris, and Jullien (1991). For some intuition on their interplay, consider again the example with normal information. The reason why learning can fail for arbitrary preferences is that society's beliefs may get stuck at some belief at which agents are taking an incorrect action, but only a strong signal about state 2 can change the action—alas no such signal is available. However, under SCD preferences, if a strong signal about state 2 would change the action, then so would a signal that indicates the state sufficiently is unlikely to be 1 or 3. Such signals are guaranteed in the normal example, and more generally by DUB information.[2]

Theorem 1 further establishes that DUB information and SCD preferences are a minimal pair of sufficient conditions of adequate learning in the sense of Athey (2002): if one property fails, then there are primitives satisfying the other property such that *in*adequate learning obtains. The scope of these necessity aspects is probed by Proposition 1 and Proposition 2 in Section 4.

So far we have highlighted for motivation the fact that normal information violates unbounded (and bounded) beliefs when there are multiple (i.e., more than two) states. But there is nothing very special in this regard about normal information. Under full support (i.e., no signal excludes any state), any information structure satisfying the widely-used monotone likelihood ratio property (MLRP) precludes unbounded beliefs with multiple states; see Proposition 4 in Section 5. Accordingly, it may be too much to ask for an information structure to guarantee learning for all preferences with multiple states. As already

[2] With a finite number of states, DUB implies that one can become sure about extreme states. But the more important property, which generalizes to an infinite number of states, is that DUB information allows one to rule *out* all states below (or above) any given state, rather than to become sure about any state. We expand on this point in Subsection 4.1.

noted, SCD is a prevalent economic restriction on preferences for ordered state spaces. DUB—which is substantially more permissive than unbounded beliefs when there are multiple states, in particular compatible with MLRP—is thus a useful informational condition that demarcates when there is learning.

To examine how broadly DUB comports with familiar signal structures, we study in Section 5 location-family signal structures. Intuitively, since DUB requires that signals be able to provide near-certainty about any state relative to its lower and upper sets, the tails of the signal distributions are key. Proposition 3 establishes that thin tails in the sense of being subexponential are sufficient for DUB; on the other hand, thick tails in the sense of being superexponential will tend to violate DUB. In conjunction with Theorem 1, it follows that location families with thin tails assure adequate learning under any SCD preferences, whereas thick tails will lead to inadequate learning for some SCD preferences.

Other related literature. Within the literature on the standard sequential Bayesian observational or social learning model, the assumption of binary states is predominant. We are aware of the following exceptions. Lee (1993) obtains learning using an infinite action space with preferences that make the optimal action sufficiently sensitive to an agent's posterior; see also Ali (2018). The logic behind learning in our framework is fundamentally different; in particular, the action space can be finite, as in most of the literature. Smith and Sørensen (2000, Section 6.1) and Arieli and Mueller-Frank (2019, Theorem 1) obtain learning for all preferences using unbounded beliefs/signals. Goeree, Palfrey, and Rogers (2006) obtain learning for any signal structure satisfying some mild conditions using "unbounded" preference shocks. To our knowledge, our paper is the first to seek a weaker informational condition that guarantees learning for homogenous agents whose preference is in a broad but circumscribed class.[3] Intimately related, the significance of an ordered state space—which we use to identify the SCD preference class—has not received attention in this literature.

Although not formally addressed by this paper, we conjecture that the interplay between SCD preferences and DUB information will be useful in moving from two to multiple states in other Bayesian observational learning settings, such as when agents only partially observe the history (Çelen and Kariv, 2004; Acemoglu, Dahleh, Lobel, and Ozdaglar, 2011; Lobel and Sadler, 2015; Smith, Sørensen, and Tian, 2021). We also leave to future research issues of how fast learning occurs under SCD preferences and DUB information. For bi-

[3] Arieli and Mueller-Frank's (2019) Theorem 3 tackles a special utility function; we elaborate in Section 5. We also note that there are papers that study Bayesian observational learning with more than two states, but they depart from the canonical setting in some important respect. For example, Gale and Kariv (2003) have agents taking decisions repeatedly, while Park and Sabourian (2011)'s financial-markets model is not one with only an information externality.

nary states, 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 observational learning, surveyed by Golub and Sadler (2016). There has also been recent interest in (mis)learning among misspecified Bayesian agents; see, for example, Frick, Iijima, and Ishii (2020) and Bohren and Hauser (2021).

2. Model

There is a countable (finite or infinite) set of ordered states Ω. We assume $\Omega \subseteq \mathbb{Z}$. There is a set of signals S; for technical simplicity, we take $S \subseteq \mathbb{R}$, either countable or an interval. There is a universal set of actions $\mathcal{A}$, with agents facing a choice from some countable subset $A \subseteq \mathcal{A}$; to avoid trivialities, any choice set has $|A| > 1$.[4]

An information or signal structure is given by a collection of cumulative distribution functions over signals, one for each state. When S is countable, let $f(s|\omega)$ denote the associated probability mass functions; when S is an interval, we assume there are probability density functions $f(s|\omega)$ such that, for convenience, for all s, $f(s|\cdot)$ is bounded. We assume no signal can exclude any state, i.e., for all ω and s, $f(s|\omega) > 0$. Note, however, that any state (or set of states) may be arbitrarily close to excluded by some signal.

The game. At the outset, date 0, a state ω is drawn from a known prior probability mass function $\mu_0 \in \Delta\Omega$. An infinite sequence of agents, indexed $n = 1, 2, \ldots$, then sequentially select actions. An agent n chooses her action $a_n \in A$ at date n, observing both a private signal s_n drawn independently from $f(\cdot|\omega)$ and the history of predecessors' actions, $h^n \equiv (a_1, \ldots, a_{n-1}) \in A^{n-1}$. No agent observes either the state or her predecessors' signals. An agent n's mixed strategy is thus a (measurable) function $\sigma_n : S \times A^{n-1} \to \Delta A$. All agents are expected utility maximizers and share the same preferences over their actions: every agent n's state-dependent preferences are represented by the von-Neumann-Morgenstern

[4] All our results except Proposition 2 go through for a general countable and ordered Ω. (By "ordered", we mean "totally ordered".) Proposition 2 uses the additional property that for every $\omega \in \Omega$, (i) if there is any strictly lower state $\omega' < \omega$, then $\max\{\omega' : \omega' < \omega\} \neq \emptyset$, and (ii) analogously if there is a strictly higher state.
The key simplicity in assuming $S \subseteq \mathbb{R}$ concerns measurability matters. Our main result, Theorem 1, does not require an order on the set of signals. But other results do presume an ordered set of signals.
Lastly, we distinguish between the choice set A and the universal set $\mathcal{A}$ because we will establish the necessity of preferences with single-crossing differences (Proposition 1 / Theorem 1 part 2) by leveraging the requirement that learning must hold for all choice sets, specifically all binary choice sets.

utility function $u : \mathcal{A} \times \Omega \to \mathbb{R}$. To ensure expected utility is well defined for all probability distributions, we assume that for all a, $u(a, \cdot)$ is bounded.

We study the Bayes Nash equilibria—hereafter simply equilibria—of this game.[5]

Adequate learning. Define the choice correspondence $c : \Delta\Omega \rightrightarrows A$ by $c(\mu) \equiv \arg\max_{a \in A} \mathbb{E}_\mu[u(a, \omega)]$. So $c(\mu)$ is the set of optimal actions under belief μ.[6] Abusing notation, we write $c(\omega) \equiv c(\delta_\omega)$, where δ_ω denotes the distribution that puts probability one on ω. Similarly to Aghion et al. (1991), we say there is *adequate knowledge* at belief μ if $\bigcap_{\omega \in \operatorname{Supp} \mu} c(\omega) \neq \emptyset$. In words, there is adequate knowledge at a belief if there is some action that is optimal at every state not precluded by that belief. Achieving adequate knowledge is a desideratum because there is no information that an agent would value obtaining given a belief if and only if that belief has adequate knowledge.

Given a prior μ_0, a signal structure f, and a strategy profile $(\sigma_n)_{n=1}^\infty$, every history $h^n \in A^{n-1}$ induces via Bayes rule a *public belief* $\mu(h^n) \in \Delta\Omega$; this is agent n's belief about ω having observed her predecessors' actions but not her private signal. Since the history is stochastic, let $\tilde{\mu}_n$ denote the corresponding random variable. By standard arguments, the stochastic process $\langle \tilde{\mu}_n \rangle$ is a martingale, and by the martingale convergence theorem, $\langle \tilde{\mu}_n \rangle \to_{a.s.} \tilde{\mu}^*$. Given a signal structure, we say there is *adequate learning* if for every prior μ_0, every choice set, and every equilibrium strategy profile, $\tilde{\mu}^*$ has adequate knowledge with probability 1.[7] In other words, there is adequate learning if no matter what prior agents begin with and what choices they are faced with, in every equilibrium almost surely agents asymptotically take correct actions. For a given signal structure f, we say there is *inadequate learning* if there is some prior and some choice set such that under every equilibrium strategy profile, there is nonzero probability that $\tilde{\mu}^*$ does not have adequate knowledge. In the natural way, we will also sometimes refer to *(in)adequate learning for a particular prior and/or choice set*.

[5] 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. Since our results concern properties of all equilibria, adopting Bayes Nash putatively strengthens our results.

[6] Throughout, we suppress the dependence of the choice correspondence c on the choice set A to reduce notation. Furthermore, we implicitly assume that for any A and μ, $c(\mu) \neq \emptyset$; standard assumptions can be invoked to guarantee that (e.g., that any A is finite).

[7] In symbols, $\Pr\left(\tilde{\mu}^* \in \{\mu : \bigcap_{\omega \in \operatorname{Supp} \mu} c(\omega) \neq \emptyset\}\right) = 1$, where the probability distribution is that induced by the prior, signal structure, and the equilibrium strategy profile.

3. Conditions on Preferences and Information

3.1. Single-Crossing Differences

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

Definition 1. Preferences represented by $u : \mathcal{A} \times \Omega \to \mathbb{R}$ 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 similar to a notion in Milgrom and Shannon (1994), but, following Kartik et al. (2019), the formulation is without an order on $\mathcal{A}$. Ignoring indifferences, SCD requires that the preference over any pair of actions can only flip once as the state changes monotonically.[8] If $\mathcal{A}$ is ordered, then SCD is implied by Milgrom and Shannon's (1994) notion or even a weaker one in Athey (2001).[9] SCD is thus a property widely satisfied by economic models; in particular, it is assured by supermodularity of u.

SCD is equivalent to the existence of a selection of optimal choices with an interval structure (cf. Kartik et al., 2019, Theorem 1.1). That is, SCD implies that for any choice set A and for any $\omega_1 < \omega_2 < \omega_3$, if action $a \in A$ is uniquely optimal at both ω_1 and ω_3 then it must also be optimal at ω_2:

$$\{a\} = c(\omega_1) = c(\omega_3) \implies a \in c(\omega_2).$$

Conversely, if SCD fails, this choice property is violated for some choice set: there are $\omega_1 < \omega_2 < \omega_3$ and a (binary) choice set A with $a \in A$ such that

$$\{a\} = c(\omega_1) = c(\omega_3) \text{ and } a \notin c(\omega_2).$$

3.2. Directionally Unbounded Beliefs

We next introduce our main condition on the signal structure.

[8] Our treatment of indifferences differs from both Milgrom and Shannon (1994) and Kartik et al. (2019) because we use a more permissive notion of single crossing. 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 et al. (2019). Our single-crossing notion is analogous to that used by Athey (2001) in the weak version of her Definition 1.

[9] SCD is, in fact, equivalent to saying that there exists some order on $\mathcal{A}$ with respect to which Athey's (2001) notion of "weak single-crossing property of incremental returns" is satisfied.

Definition 2. The signal structure has *directionally unbounded beliefs* (DUB) if for every ω, there are $C > 0$ and signal sequences $(\overline{s}_i)_{i=1}^{\infty}$ and $(\underline{s}_i)_{i=i}^{\infty}$ such that:

1. for all $\omega' < \omega$: (i) $\lim_{i\to\infty} \frac{f(\overline{s}_i|\omega')}{f(\overline{s}_i|\omega)} = 0$ and (ii) $\frac{f(\overline{s}_i|\omega')}{f(\overline{s}_i|\omega)} < C$ for all i;
2. for all $\omega' > \omega$: (i) $\lim_{i\to\infty} \frac{f(\underline{s}_i|\omega')}{f(\underline{s}_i|\omega)} = 0$ and (ii) $\frac{f(\underline{s}_i|\omega')}{f(\underline{s}_i|\omega)} < C$ for all i.

The reader should bear in mind that when S is an interval, DUB's requirements are implicitly required to hold for sets of signals of positive (Lebesgue) measure. Equivalently, the requirements as stated must hold no matter which version of the densities $f(\cdot|\omega)$ are used. Under familiar assumptions, it is sufficient to check the requirements for a given choice of the densities: e.g., if Ω is finite and each $f(\cdot|\omega)$ is continuous, or more generally, the family $\{f(\cdot|\omega)\}_{\omega\in\Omega}$ is equicontinuous. A similar point applies to some other properties below; to ease the exposition, we will suppress the positive-measure caveat hereafter.

The key idea of DUB consists of point (i) in each part of the definition; indeed, each point (ii) can be safely ignored when Ω is finite. To interpret DUB, fix any state ω. Part 1 of DUB implies that there are signals that are arbitrarily relatively more likely in ω than in all $\omega' < \omega$. We stress that while the sequence $(\overline{s}_i)$ can depend on ω, it cannot depend on the lower state $\omega' < \omega$. The idea is that starting from any prior with support $\{\omega' : \omega' \leq \omega\}$, there are signals that lead to posteriors arbitrarily near certainty about ω; in this sense, ω is (fully) distinguishable from all lower states, simultaneously.

Part 2 of DUB's definition analogously implies distinguishability of ω from all higher states, simultaneously. Such distinguishability actually characterizes DUB when Ω is finite. We make these points precise as follows.

Definition 3. State ω is *distinguishable* from states $\Omega' \subseteq \Omega \setminus \{\omega\}$ if for every $\mu \in \Delta\Omega$ with $\mu(\omega) > 0$, there is a sequence of signals $(s_i)_{i=1}^{\infty}$ such that

$$\lim_{i\to\infty} \frac{\sum_{\omega'\in\Omega'} \mu(\omega')f(s_i|\omega')}{\mu(\omega)f(s_i|\omega)} = 0. \tag{1}$$

There is *directional distinguishability* if every ω is distinguishable from $\{\omega' : \omega' < \omega\}$ and also distinguishable from $\{\omega' : \omega' > \omega\}$.

Note that the fraction on the left-hand side of Equation 1 is the posterior likelihood ratio of Ω' relative to ω given prior μ and signal s_i.

Lemma 1. *DUB implies directional distinguishability. If Ω is finite, directional distinguishability implies DUB.*

Crucially, DUB does not require or assure that any state ω is distinguishable from all other states—not even any subset of states containing both some higher and some lower state than ω. Distinguishability of each state from all other states corresponds to the multi-state version of Smith and Sørensen's (2000) unbounded beliefs; we elaborate in Section 5. What DUB does assure, and will be crucial in our analysis, is that for any ω and prior μ with $\mu(\omega) > 0$, there are signals that make the posterior on the lower set $\{\omega' : \omega' < \omega\}$ arbitrarily small, and analogously for the upper set $\{\omega' : \omega' > \omega\}$.

A leading example of DUB information is when signals are normally distributed with mean ω and fixed variance; we hereafter refer to this specification as the *normal signal structure* or just *normal information*. In this case, for any $\omega' < \omega$, the ratio $f(s|\omega')/f(s|\omega) \to 0$ as $s \to \infty$, while the same ratio tends to ∞ as $s \to -\infty$; hence, for any ω, any sequence of signals that is unbounded above (resp., below) verifies part 1 (resp., part 2) of DUB's definition.[10] Notably, however, no state can be simultaneously distinguished from both some lower one and some higher one. Figure 1a provides a graphical depiction. The solid curve therein is the set of posteriors that obtain under normal information when $\Omega = \{1, 2, 3\}$, given prior μ. Posteriors are bounded away from certainty on state 2. Yet they can become arbitrarily certain about the relative probability of state 2 to 1 (as the curve approaches state 3) and, separately, about the relative probability of state 2 to 3 (as the curve approaches state 1)—this corresponds graphically to the solid curve being tangent to the relevant edges at the two vertices.

Figure 1c depicts a failure of DUB.[11] The information structure corresponding to the triangular set of posteriors has the property that no signal can distinguish state 3 from states 2 and 1 simultaneously, even though a signal leading to a posterior close to the 1–3 edge (resp., the 2–3 edge) distinguishes state 3 from 2 (resp., from 1). Another failure of DUB is seen in Figure 2 in Subsection 4.2. There, states 1 and 3 are distinguishable from their complements, as the posteriors under the full-support prior μ', given by the black

[10] To verify point (ii) for part 1 of the definition, note that for any $\omega' < \omega$ and any $s > \omega$, normal information implies $f(s|\omega')/f(|\omega) < 1$, and hence we can choose the bound $C = 1$ with each element of the sequence $(\overline{s}_i)$ above ω. An analogous observation applies for part 2 of Definition 2. These observations are a special case of Proposition 3 that applies to location-shift signal structures.

[11] 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. Moreover, the set of posteriors for any full-support prior μ identifies the set of posteriors for any other prior μ'. Specifically, the posteriors for any prior μ' can then be computed using the following implication of Bayesian updating: denoting the posterior after signal s with a subscript s, it holds that for any $\{\omega, \omega'\} \subseteq$ Supp μ',

$$\frac{\mu'_s(\omega)}{\mu'_s(\omega')} = \frac{\mu_s(\omega)}{\mu_s(\omega')} \frac{\mu'(\omega)}{\mu'(\omega')} \frac{\mu(\omega')}{\mu(\omega)}.$$

curve, approach the relevant vertices. However, state 2 is not distinguishable from state 1: the posteriors for the prior μ with support $\{1, 2\}$, given by the red line, do not approach state 2's vertex; equivalently, the black curve is not tangent to the 2–3 edge at state 3's vertex.

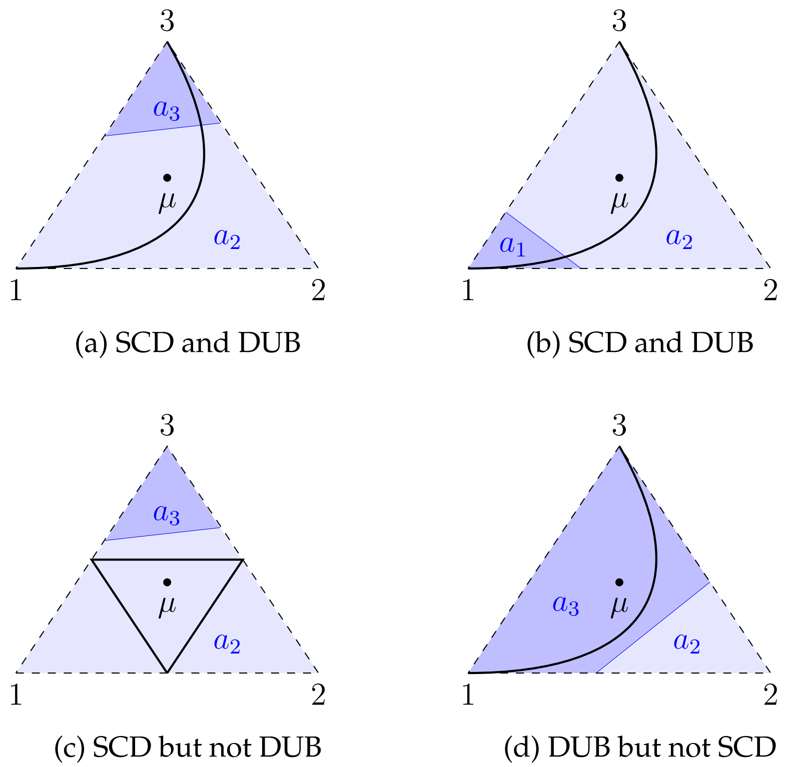

Figure 1: For belief μ, posteriors in black solid curve/triangle. Preference regions among actions shaded in blues.

We should note that when the state space Ω is infinite, Lemma 1's conclusion that DUB implies directional distinguishability relies on point (ii) of each part of DUB's definition, which we refer to as a *uniform boundedness* requirement on the likelihood ratios. Example 2 in Appendix B fleshes the point out, but the gist is as follows. Consider $\Omega = \{\dots, -1, 0\}$. There are signal structures such that for any given $\omega' < 0$, $f(s|\omega')/f(s|0) \to 0$ as $s \to \infty$, but $\lim_{s\to\infty} \sup_{\omega'<0} f(s|\omega')/f(s|0) = \infty$. The latter point means that any unbounded-above sequence of signals violates uniform boundedness. While state 0 is distinguishable from any finite set of lower states, we show that it is not distinguishable from all lower states. Specifically, there are full-support priors such that given any state $\omega' < 0$, arbitrarily high signals make it arbitrarily more likely that the state is 0 than ω', while at the same time making it arbitrarily close to certain that the state is not 0.

Although DUB's uniform boundedness requirement assures directional distinguishability, the requirement may not be necessary for directional distinguishability when Ω is infinite (cf. Lemma 1). Appendix B elaborates. We establish there that for an infinite Ω DUB characterizes a strengthening of directional distinguishability that imposes prior independence, viz., the version of directional distinguishability in which the sequence of signals in Definition 2 cannot depend on the prior. Moreover, we establish that under a standard informational condition, prior-independent directional distinguishability—and hence DUB—is equivalent to directional distinguishability.

4. Main Results

Our main result is that SCD preferences and DUB information combine to ensure adequate learning; moreover they are a minimal pair of sufficient conditions, at least when Ω is finite.

Theorem 1.

1. *If preferences have SCD and the signal structure has DUB, then there is adequate learning.*
2. *If preferences fail SCD, then there is a signal structure with DUB for which there is inadequate learning.*
3. *If the signal structure fails DUB and Ω is finite, then there are preferences with SCD for which there is inadequate learning.*

(All formal proofs not in the main text are in Appendix A.)

To explain the logic behind Theorem 1, we first describe how adequate learning can be reduced to checking whether certain beliefs have adequate knowledge. Given some preferences, signal structure, and choice set, say that a (public) belief μ is *stationary* if there is some optimal action at μ that almost surely remains optimal no matter the signal an agent (privately) observes. That is, stationary beliefs are those at which information from the signal structure at hand has no further value, or equivalently, those at which there can be an information cascade.[12] A necessary condition for adequate learning is that all stationary beliefs have adequate knowledge: if a stationary belief μ has inadequate knowledge, then when the prior is μ there is an equilibrium with an immediate herd on an action that is suboptimal in some state in μ's support. In fact, it can be shown that if all stationary beliefs have adequate knowledge, then there is adequate learning (for the given choice set). The reason is that all limit public beliefs must be stationary; see Arieli and Mueller-Frank (2019,

[12] There is an information cascade at a history if, almost surely, every subsequent agent's action is independent of her private signal. There is a herd at a history if all subsequent agents take the same action. Modulo tie-breaking issues, an information cascade implies a herd but not vice-versa.

Lemma 2). For, any nonstationary public belief will lead to an agent taking different actions after different signals, and this will generate new public beliefs that are bounded away from the original one. Accordingly, the question of whether there is adequate learning is equivalent to the question of whether every stationary belief has adequate knowledge.

4.1. Sufficiency

In general, stationary beliefs can have inadequate knowledge because of a mismatch between preferences and information: there may be no signals that provide enough information about the states that agents would benefit from learning about.

However, part 1 of Theorem 1 holds because SCD and DUB together preclude such a mismatch. If μ has inadequate knowledge, then $c(\omega^*) \neq c(\mu)$ for some state $\omega^* \in \operatorname{Supp} \mu$. (For simplicitly, suppose $c(\mu)$ and $c(\omega^*)$ are singletons.) SCD implies that not only is $c(\omega^*)$ strictly preferred to $c(\mu)$ at ω^*, but it is also weakly preferred either at all higher states or at all lower states.[13] Suppose the former case. The directional distinguishability implied by DUB ensures that the posterior likelihood ratio of ω^* to its lower set $\{\omega' : \omega' < \omega^*\}$ can be arbitrarily high. This implies the posterior on $\{\omega' : \omega' \geq \omega^*\}$ can be arbitrarily close to one. Hence, with nonzero probability, an agent will strictly prefer $c(\omega^*)$ to $c(\mu)$. So μ is not stationary. Consequently, under SCD and DUB, all stationary beliefs have adequate knowledge. Figure 1a and 1b illustrate the idea: since action a_2 is optimal at μ whose support is $\{1, 2, 3\}$, if μ does not have adequate knowledge then either a different action must be optimal at state 3 (panel 1a) or at state 1 (panel 1b). In both cases there is nonzero probability of the posterior falling in the relevant region (corresponding to action a_3 and a_1, respectively).

We would like to caution against the following perspective on how SCD and DUB work together to assure learning. With a finite Ω one might focus on extreme states, as in Figure 1a and 1b: given any nondegenerate belief with inadequate knowledge, SCD implies distinct optimal actions at the belief's extreme states (ignoring indifferences); DUB implies these extreme states are distinguishable from all others in the belief's support; hence the belief is not stationary.[14] But the notion of extreme states is ill-defined for beliefs with infi-

[13] For a given choice set, it is this property that is key. Since we allow for arbitrary choice sets, it is equivalent to SCD. With restrictions on the allowable choice sets, weaker properties than SCD may suffice. For example, if the choice set must be an interval (with respect to some exogenous order on $\mathcal{A}$), then Quah and Strulovici's (2009) interval dominance order is sufficient.

[14] The extreme states of a belief μ with finite support are $\max\{\omega : \mu(\omega) > 0\}$ and $\min\{\omega : \mu(\omega) > 0\}$. Note that we must consider beliefs with arbitrary support because the limit public belief need not have full support, even though the public belief at every finite time does. Both SCD and DUB ensure their choice and distinguishability properties, respectively, at the extreme states of any belief, not just at the extreme states of full-support beliefs.

nite support, and in that case DUB is compatible with no state being distinguishable from all others. Normal information is a case in point. So nothing about extreme states can be essential to assure learning. Rather, DUB's directional distinguishability implies that given any belief and any state in the support, signals will be able to (almost) rule out all lower states and, separately, all higher states. SCD ensures that such distinguishability is sufficient by the logic explained in the previous paragraph, not simply because of distinct optimal actions in certain states. Indeed, although pairing DUB with distinct optimal actions in all states (or, more generally, just interval choice) for a given choice set would be sufficient for learning for that choice set when there is only a finite number of states, the following example shows that it does not guarantee learning in general.

Example 1. $\Omega = \mathbb{Z}$, there is normal information, and the choice set is $A = \mathbb{Z} \cup \{a^*\}$ with preferences in any state ω given by: $u(\omega, \omega) = \varepsilon$; $u(a^*, \omega) = 0$; and $u(a, \omega) = -1/\varepsilon$ for $a \notin \{a^*, \omega\}$. The parameter $\varepsilon > 0$ is small. So in any state ω the uniquely optimal action is ω, but this action gives very low utility in any other state. Action a^* is a uniformly safe action. Note that SCD is violated. There are priors such that the posterior probability of any state is uniformly bounded away from one (no matter the signal or state),[15] and hence for small enough $\varepsilon > 0$, the safe action a^* is always optimal. Any such prior is stationary, even though it has inadequate knowledge. $\square$

4.2. Necessity

Necessity of SCD. To understand part 2 of Theorem 1, assume a violation of SCD. That is, there are a pair of actions and three states such that one action is strictly preferred at the low and high states, while the other is strictly preferred at the middle state. See Figure 1d: a_3 is strictly preferred to a_2 in states 1 and 3 while the preference is reversed in state 2. The signal structure depicted there has DUB; it depicts normal information. Plainly, belief μ, which has inadequate knowledge, is stationary when agents only choose between actions a_2 and a_3.[16]

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

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

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

[16] The logic of Figure 1d implies that if any state is not distinguishable from its complement, then (for any choice set) there are non-SCD preferences and some stationary belief with inadequate knowledge.

Theorem 1 part 2 can actually be strengthened substantially. Recall the *monotone likelihood ratio property (MLRP)*: for every $\omega' > \omega$ and $s' > s$, $\frac{f(s'|\omega')}{f(s'|\omega)} \geq \frac{f(s|\omega')}{f(s|\omega)}$ (see, e.g., Milgrom, 1981). MLRP is a strong but widely-used property satisfied by many familiar signal structures, including normal information.

Proposition 1. *If preferences fail SCD, then for any signal structure with MLRP there is inadequate learning.*

Here is the logic. Consider $\Omega = \{1, 2, 3\}$. MLRP implies that if a signal s provides a higher likelihood of 2 relative to 1 than does signal s', then s provides a lower likelihood of state 2 relative to 3 than s'. Hence, there is no signal that simultaneously provides near-certainty about state 2 relative to both 1 and 3. A violation of SCD means that there is choice set $\{a, a'\}$ with action a strictly preferred to a' at state 2 but a' strictly preferred at both 1 and 3. Any prior that puts sufficiently low (but nonzero) probability on state 2 relative to both 1 and 3 is stationary, even though there is inadequate knowledge.

Proposition 1 and Theorem 1 part 2 use the fact that adequate learning requires adequate learning for all choice sets (and priors). For some choice sets, adequate learning may be assured even without SCD; after all, a violation of SCD may be among dominated actions that are irrelevant for that choice set. That said, if for a given choice set there are states $\omega_1 < \omega_2 < \omega_3$ such that $\{a_3\} = c(\omega_1) = c(\omega_3)$ and $\{a_2\} = c(\omega_2)$, then any MLRP signal structure has inadequate learning for that choice set by the logic in the previous paragraph.

Necessity of DUB. Turning to part 3 of Theorem 1, a failure of DUB when Ω is finite means that some state ω^* cannot be distinguished from either all lower states or all upper states (Lemma 1).[17] Suppose it is the former case. For any two actions a^* and $\underline{a}$, there are SCD preferences such that a^* is strictly preferred in state ω^* while $\underline{a}$ is strictly preferred in all lower states. This can be shown to imply that any belief with support $\{\omega : \omega \leq \omega^*\}$ that puts sufficiently low probability on ω^* relative to every $\omega < \omega^*$ is stationary. Since any such belief does not have adequate knowledge, there is inadequate learning. See Figure 1c for the idea, where state 3 cannot be distinguished from $\{1, 2\}$, but μ is stationary because all induced posteriors lead to action a_2.

Learning at a Fixed Prior. One may also be interested in whether DUB is necessary for adequate learning (presuming SCD) at a given prior, in particular at an arbitrary full-support

[17] When Ω is infinite, the ensuing logic goes through with a failure of directional distinguishability (which is no longer assured by a failure of DUB).

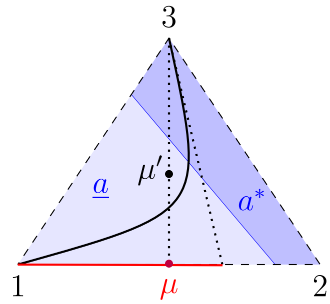

Figure 2: Preference regions among the actions $\underline{a}$ and a^* shaded in blues. An MLRP signal structure such that under belief μ', posteriors are given by the black curve, while under belief μ, posteriors are given by the red line.

prior. To see why the prior's support is relevant, consider Figure 2. DUB fails because state 2 is not distinguishable from state 1. Action a^* is optimal at states 2 and 3 while $\underline{a}$ is optimal at state 1. On the one hand, if we consider all priors, it is straightforward that there is inadequate learning: belief μ, which is supported on $\{1, 2\}$, has inadequate knowledge and is stationary. On the other hand, because the optimal actions are distinct at the extreme states 1 and 3, and the extreme states are distinguishable from their complements, no full-support belief is stationary (this contrasts with Figure 1c). Since at any finite time the public belief's support is always that of the prior, the public belief at any finite time in this example is nonstationary, i.e., there cannot be a finite-time information cascade starting from any full-support prior. This observation may raise a concern about establishing the necessity of DUB in part 3 of Theorem 1 using priors with restricted support. However:

Proposition 2. *If the signal structure fails DUB but satisfies MLRP, then there are SCD preferences such that for any full-support prior there is inadequate learning.*

In particular, the signal structure and preferences in Figure 2 entail inadequate learning for any full-support prior, such as μ' in the figure.[18] We highlight the implication that, even under SCD, adequate learning can fail for every full-support prior if DUB is weakened to only ensuring that all upper and lower sets of states are distinguishable from their complements.

Here is the key idea behind Proposition 2. Assume MLRP. We will argue that if for every SCD preference there is a full-support prior for which there is adequate learning, then directional distinguishability holds. Akin to the explanation earlier in this subsection, pick any two actions a^* and a', any state ω^*, and consider SCD preferences such that a' is strictly preferred to a^* at every state $\omega < \omega^*$ while a^* is strictly preferrred at every $\omega \geq \omega^*$ (and

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

all other actions are dominated, hence can be ignored). Assume there is adequate learning at some full-support prior. MLRP then guarantees the following property: with positive probability there is an infinite history with a herd on a', call it $h' \equiv (\underbrace{a_1, \ldots, a_n}_{\text{each } a_i \in \{a', a^*\}}, a', a', \ldots)$, whose limit belief μ' satisfies $\operatorname{Supp} \mu' = \{\omega : \omega < \omega^*\}$.[19] This support property can be restated as $\Pr(h'|\omega) > 0 \iff \omega < \omega^*$. So the probability of signals that overturn the herd—i.e., lead to action a^*—must vanish over time at a fast enough rate in each $\omega < \omega^*$, but either not vanish or vanish at a slow enough rate in each $\omega \geq \omega^*$. In particular, there must exist overturning signals whose probability gets arbitrarily large in state ω^* relative to those in every $\omega < \omega^*$, which implies that ω^* is distinguishable from its lower set $\{\omega : \omega < \omega^*\}$. A symmetric argument using preferences in which a^* is optimal in state ω^* and below while a' is optimal in states above ω^* establishes that ω^* is distinguishable from its upper set $\{\omega : \omega > \omega^*\}$. Since ω^* was arbitrary, it follows that there is directional distinguishability. Finally, under MLRP, directional distinguishability implies DUB (Corollary 1 in Appendix B).

The foregoing argument reveals that MLRP's role in Proposition 2 is unexpected. As elaborated in fn. 19, MLRP allows us to establish that for the preferences specified in the previous paragraph there is a positive-probability limit belief whose support contains *all* states in which action a' is optimal. We suspect that MLRP is not needed to guarantee this property. To see why, note that the preferences could be chosen so that all states in which a' is optimal are utility equivalent. (A set of states is utility equivalent if for every action a, $u(a, \cdot)$ is constant over that set.) Say that arbitrary preferences and information satisfy *non-exclusion of equivalent states* if for every prior and every set of utility-equivalent states Ω' there is an equilibrium in which there is a limit belief $\mu' \in \operatorname{Supp} \tilde{\mu}^\infty$ such that $\Omega' \subseteq \operatorname{Supp} \mu'$. While we do not have a proof, it seems plausible that non-exclusion of equivalent states would hold generally when there is adequate learning (for any preferences and information)—otherwise, in all equilibria, there are herds that rule out some state in which the action being herded on is optimal. If non-exclusion of equivalent states is assured when there is adequate learning, the argument establishing directional distinguishability in the

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

previous paragraph goes through without invoking MLRP.[20]

5. Discussion

This section elaborates on aspects of DUB, including its relationship with earlier notions of unbounded beliefs.

5.1. DUB in Location Families

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

Formally, we say that the signal structure is a *location-shift signal structure* if $S = \mathbb{R}$ and there is a density $g : \mathbb{R} \to \mathbb{R}_+$, referred to as the *standard density*, such that $f(s|\omega) = g(s-\omega)$. For simplicity, we restrict attention to standard densities that are bounded. We say that g is *strictly subexponential* if there is some $p > 1$ ($p \in \mathbb{R}$) such that for all s large enough in absolute value, $g(s) \leq \exp(-|s|^p)$. Normal information satisfies this condition with any $p \in (1, 2)$.[21] Intuitively, a strictly subexponential density has a thin tail—it eventually decays faster than the exponential density.

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

The exponent p being strictly larger than one in the definition of strictly subexponential is essential for the result. To see that, consider the Laplace or doublexponential standard density $g(s) = (1/2)\exp(-|s|)$. This density is not strictly subexponential, and indeed DUB fails: for any $\omega' < \omega < s$, $f(s|\omega')/f(s|\omega) = g(s-\omega')/g(s-\omega) = \exp(\omega' - \omega)$ is independent of s.[22] More broadly, thick-tailed standard densities will tend to violate DUB, intuitively because a thick tail implies that fixing any pair of states, extreme signals become

[20] To see where the previous paragraph's argument for directional distinguishability fails without the condition, suppose that all limit beliefs are perfectly informative (i.e., have singleton supports). Then any infinite history with a herd on a' can have positive probability only in a particular state $\omega < \omega^*$. Overturning this herd in state ω^* would then only imply that ω^* is distinguishable from ω, rather than ω^* being distinguishable from its lower set.

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

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

uninformative.[23] For example, DUB fails for standardized student t-distribution with arbitrary degrees of freedom: it can be checked that for any $\omega' < \omega$, as $s \to \infty$ or $s \to -\infty$, $f(s|\omega')/f(s|\omega) \to 1$. By Theorem 1 part 3, such signal structures will entail inadequate learning under some SCD preferences.

5.2. Other Informational Properties

It is instructive to compare DUB with some other notions of unbounded beliefs. For simplicity, we restrict attention hereafter to a finite set of states, in which case the uniform boundedness requirement of DUB is automatic.

Definition 4. The signal structure has:

1. *universal directionally unbounded beliefs* if there are signal sequences $(\overline{s}_i)_{i=1}^{\infty}$ and $(\underline{s}_i)_{i=i}^{\infty}$ such that for every ω: (i) for all $\omega' < \omega$, $\lim_{i\to\infty} \frac{f(\overline{s}_i|\omega')}{f(\overline{s}_i|\omega)} = 0$, and (ii) for all $\omega' > \omega$, $\lim_{i\to\infty} \frac{f(\underline{s}_i|\omega')}{f(\underline{s}_i|\omega)} = 0$.
2. *pairwise unbounded beliefs* if for every ω and $\omega' \neq \omega$, there is a sequence $(s_i)_{i=1}^{\infty}$ such that $\lim_{i\to\infty} \frac{f(s_i|\omega')}{f(s_i|\omega)} = 0$.
3. *unbounded beliefs* if for every ω there is a sequence $(s_i)_{i=1}^{\infty}$ such that for all $\omega' \neq \omega$: $\lim_{i\to\infty} \frac{f(s_i|\omega')}{f(s_i|\omega)} = 0$.

Universal DUB strengthens DUB by not allowing the sequences $(\overline{s}_i)$ and $(\underline{s}_i)$ that verify DUB to depend on the benchmark state ω. Following the discussion in Subsection 3.2, it is clear that normal information satisfies universal DUB. On the other hand, pairwise unbounded beliefs, introduced by Arieli and Mueller-Frank (2019), is weaker than DUB because it allows the sequences $(\overline{s}_i)$ and $(\underline{s}_i)$ to depend on not only the benchmark state ω but also the comparison state ω'. Put differently, pairwise unbounded beliefs requires only that given any pair of states, each state be distinguishable from the other. Unbounded beliefs, following Smith and Sørensen (2000), is stronger than DUB by requiring that the two sequences $(\overline{s}_i)$ and $(\underline{s}_i)$ that verify DUB be, effectively, the same. Put differently, unbounded beliefs requires that every state be distinguishable from its complement.

With only two states, all four versions of unbounded beliefs—unbounded beliefs, universal DUB, DUB, and pairwise unbounded beliefs—are equivalent. But more generally they are not.

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

Proposition 4. *The following relationships hold:*

1. *Unbounded beliefs $\implies$ DUB $\implies$ pairwise unbounded beliefs.*
2. *Universal DUB $\implies$ DUB $\implies$ pairwise unbounded beliefs.*
3. *If MLRP holds, then pairwise unbounded beliefs $\iff$ DUB $\iff$ universal DUB.*
4. *If $|\Omega| > 2$, then MLRP $\implies$ not unbounded beliefs.*

Proof of Proposition 4. Parts 1 and 2 are clear from the definitions. For part 3, it suffices to show that MLRP and pairwise unbounded beliefs imply universal DUB. MLRP and pairwise unbounded beliefs imply 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, requirement (i) of universal DUB is satisfied by any sequence $(s_i) \to \sup S$ and requirement (ii) by any sequence $(s_i) \to \inf S$. For part 4, assume $\omega_1 < \omega_2 < \omega_3$ and MLRP. Since $f(s|\omega_2)/f(s|\omega_3)$ is decreasing in s while $f(s|\omega_2)/f(s|\omega_1)$ is increasing, there is no sequence of signals for which both these ratios can tend to 0; hence unbounded beliefs fails. *Q.E.D.*

The relationships in parts 1 and 2 of Proposition 4 cannot be strengthened. Figure 1c shows that pairwise unbounded beliefs does not imply DUB. As we have already noted, normal information (which satisfies universal DUB and MLRP) does not have unbounded beliefs when there are more than two states. Hence, it illustrates that universal DUB does not imply unbounded beliefs. Figure 3a shows that unbounded beliefs does not imply universal DUB either: there is no single sequence of signals that distinguishes state 3 from its lower states while also distinguishing state 2 from its lower state. Figure 3b shows, however, that univeral DUB and unbounded beliefs are compatible—simply note that all posteriors are possible. Figure 3c shows that DUB does not imply universal DUB: universal DUB does not hold for the same reason as in Figure 3a; to confirm DUB, notice that state 2 can be distinguished from state 1 (resp., 3) by signals approaching the interior of the 2–3 (resp., 1–2) edge. Observe that, consistent with parts 3 and 4 of Proposition 4, none of the panels in Figure 3 have MLRP: in each panel there is a sequence of signals that makes state 2 more likely than both state 1 and state 3 simultaneously.

Parts 3 and 4 of Proposition 4 are notable because of MLRP's widespread application in information economics. Part 4 is a manifestation of the restrictiveness of unbounded beliefs with more than two states. We have seen that unbounded beliefs is not necessary for adequate learning for SCD preferences. On the other hand, pairwise unbounded beliefs is necessary for adequate learning (in a sense analogous to Theorem 1 part 2) given any

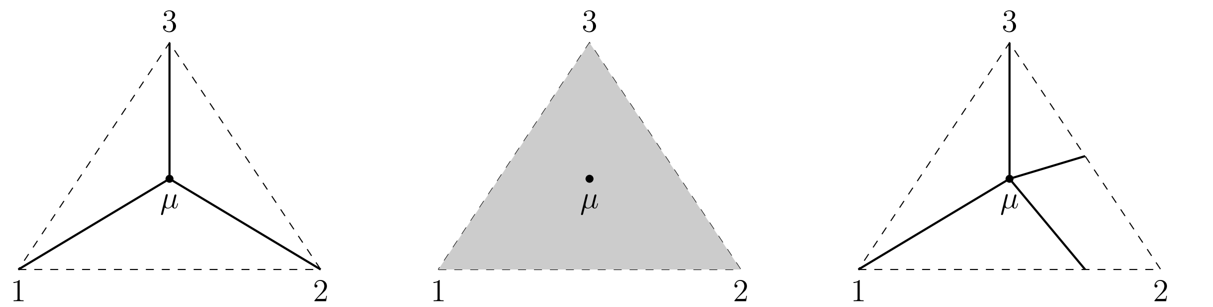

(a) Unbounded beliefs but not universal DUB

(b) Unbounded beliefs and universal DUB

(c) DUB but neither unbounded beliefs nor universal DUB

Figure 3: Prior μ. Posteriors in shaded areas/solid lines.

minimally-rich class of preferences: so long as for any pair of states, there is a utility function representing different preferences in these states.[24] Part 3 of Proposition 4 thus implies that under MLRP, universal DUB is necessary for adequate learning for any minimally-rich class of preferences.

5.3. Existing Results

It is straightforward to see that unbounded beliefs is necessary to assure adequate learning for all utility functions and priors; recall fn. 16. Smith and Sørensen (2000) establish that it is also sufficient—formally for two states, but as they say in their Section 6.1, the logic applies with any finite number of states.[25] Arieli and Mueller-Frank (2019, Theorem 1) extend these points to a model with more general signal, state, and action spaces, referring to unbounded beliefs as "totally unbounded". Plainly, any order on Ω becomes irrelevant once one allows for arbitrary preferences. It is to be expected that the resulting condition on information—unbounded beliefs—is also independent of any order on Ω. As we have noted, however, unbounded beliefs is extremely demanding with multiple states. Given the breadth and economic relevance of SCD preferences, which presumes an order on Ω, we view it as natural to study learning for preferences in that class. We have shown that the resulting informational demand for adequate learning substantially weakens from unbounded beliefs to DUB.

[24] Suppose pairwise unbounded beliefs does not hold. Then some state ω cannot be distinguished from another state ω'. Consider strict preferences that are distinct in these two states: there are a and a' such that $u(a,\omega) > u(a',\omega)$ and $u(a,\omega') < u(a',\omega')$. Then, when the choice set is $A = \{a, a'\}$, $c(\omega) \cap c(\omega') = \emptyset$ and hence a prior with support $\{\omega, \omega'\}$ assigning probability close to 1 to ω' is stationary but does not have adequate knowledge.

[25] On the other hand, bounded beliefs ensures inadequate learning for all nontrivial preferences (i.e., there is no action that is optimal in all states). The correct notion of bounded beliefs here is that for every pair of states $\omega \neq \omega'$, $f(s|\omega)/f(s|\omega')$ is bounded above in s. That is, no state is distinguishable from any other.

Arieli and Mueller-Frank's (2019) Theorem 3 concerns pairwise unbounded beliefs. They establish that this condition ensures adequate learning for a special utility function: $A = \Omega$ (finite) and $u(a, \omega) = \mathbb{1}\{a = \omega\}$. Although this utility function satisfies SCD, the reason pairwise unbounded beliefs (rather than DUB) is sufficient here has to do with a particular property of these preferences that is independent of any ordering of the states. Specifically, these preferences induce choices that are one-to-one with the most likely state (ignoring ties). Since pairwise unbounded beliefs implies that under any nondegenerate belief, the most-likely state can be overturned (i.e., with nonzero probability, $\arg\max_{\omega \in \Omega} \mu(\omega)$ will be disjoint for the prior and the posterior), only degenerate beliefs—which are equivalent to those with adequate knowledge here—are stationary.

An interesting by-product of our results is that appending MLRP to pairwise unbounded beliefs expands adequate learning from the above special preference to all SCD preferences, simply because MLRP and pairwise unbounded beliefs imply DUB (Proposition 4).

Arieli and Mueller-Frank's (2019) Theorem 3 further establishes that for the particular utility function studied there, pairwise unbounded beliefs is also necessary for adequate learning at a given prior. Since under MLRP pairwise unbounded beliefs is equivalent to DUB, their argument provides an alternative proof of our Proposition 2. However, as explained in our discussion after the proposition, it is the more general property of non-exclusion of equivalent states that really drives the necessity of DUB for adequate learning under SCD preferences.

A. Omitted Proofs

Proof of Lemma 1. We first prove that DUB implies directional distinguishability. (It is dealing with the possibility of an infinite Ω that requires a non-immediate argument.) Consider any ω and fix both $C > 0$ and the sequence $(\overline{s}_i)_{i=1}^{\infty}$ from part 1 of DUB's definition. Take any $\varepsilon > 0$ and any belief $\mu \in \Delta\Omega$ with $\mu(\omega) > 0$. For each integer $k > 0$, define

$$M_k \equiv \{\omega' : \omega' < \omega \text{ and } \mu(\omega') < 1/k\}.$$

Let $\mu(\Omega') \equiv \sum_{\omega' \in \Omega'} \mu(\omega')$ for any set $\Omega' \subseteq \Omega$. Since $M_{k+1} \subseteq M_k$, we have $\lim_{k\to\infty} \mu(M_k) = \mu(\bigcap_{k=1}^{\infty} M_k) = 0$, and for all large k, $\mu(M_k) < \frac{\varepsilon\mu(\omega)}{2C}$. Moreover, $\overline{M}_k \equiv \{\omega' : \omega' < \omega\} \setminus M_k$ is finite, and so there exists N such that for every $i \geq N$, $\frac{f(\overline{s}_i|\omega')}{f(\overline{s}_i|\omega)} < \frac{\varepsilon\mu(\omega)}{2\sum_{\omega'' \in \overline{M}_k} \mu(\omega'')}$ for all $\omega' \in \overline{M}_k$. It follows that for every $i \geq N$,

$$\begin{aligned}\frac{\sum_{\omega'<\omega} \mu(\omega')f(\overline{s}_i|\omega')}{\mu(\omega)f(\overline{s}_i|\omega)} &= \frac{\sum_{\omega' \in M_k} \mu(\omega')f(\overline{s}_i|\omega')}{\mu(\omega)f(\overline{s}_i|\omega)} + \frac{\sum_{\omega' \in \overline{M}_k} \mu(\omega')f(\overline{s}_i|\omega')}{\mu(\omega)f(\overline{s}_i|\omega)} \\ &\leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.\end{aligned}$$

Hence, ω is distinguishable from $\{\omega' : \omega' < \omega\}$. An analogous argument establishes distinguishability of ω from $\{\omega' : \omega' > \omega\}$.

We next prove that assuming Ω is finite, directional distinguishability implies DUB. Consider any ω and the belief μ uniformly distributed over $\{\omega' : \omega' \leq \omega\}$. The distinguishability of ω from $\{\omega' : \omega' < \omega\}$ implies that for every $\varepsilon > 0$ there is a signal $\overline{s}$ such that $\frac{\sum_{\omega'<\omega} f(\overline{s}|\omega')}{f(\overline{s}|\omega)} < \varepsilon$, and so $\frac{f(\overline{s}|\omega')}{f(\overline{s}|\omega)} < \varepsilon$ for every $\omega' < \omega$. Analogously, for every $\varepsilon > 0$ there is a signal $\underline{s}$ such that $\frac{f(\underline{s}|\omega')}{f(\underline{s}|\omega)} < \varepsilon$ for every $\omega' > \omega$. It follows that DUB holds. *Q.E.D.*

Proof of Theorem 1 part 1. Fix any choice set, prior, SCD preferences, and DUB information. Let μ_s denote the posterior given public belief μ and signal s. A belief μ is *stationary* if there is some $a \in c(\mu)$ such that $a \in c(\mu_s)$ for all s, except possibly for a set of signals of zero (Lebesgue) measure when S is an interval. In other words, at a stationary public belief, the value of a private signal is zero.

Arieli and Mueller-Frank's (2019) Lemma 2 establishes that in any equilibrium, the asymptotic public belief is stationary with probability one. It therefore suffices for us to show that SCD and DUB jointly ensure that any stationary belief has adequate knowledge.

Take any belief μ and any optimal action $a \in c(\mu)$. If μ has inadequate knowledge, there exists $\omega^* \in \operatorname{Supp}\mu$ such that $a \notin c(\omega^*)$, i.e., there is an action a' such that $D_{a',a}(\omega^*) \equiv$

$u(a',\omega^*) - u(a,\omega^*) > 0$. Since $a \in c(\mu)$, there is some ω' such that $D_{a',a}(\omega') < 0$. Assume $\omega' < \omega^*$. (The proof for $\omega' > \omega^*$ is symmetric.) Since $D_{a',a}(\cdot)$ is single crossing, $D_{a',a}(\omega) \geq 0$ for every $\omega > \omega^*$. Hence, denoting $\overline{D} \equiv \sup_{\omega\in\Omega} |D_{a',a}(\omega)| < \infty$ (the inequality owes to our maintained assumption of bounded utility), it holds for any signal s that

$$\begin{aligned}
\mathbb{E}_{\mu_s}[D_{a',a}(\omega)] &\geq \sum_{\omega<\omega^*} D_{a',a}(\omega)\mu_s(\omega) + D_{a',a}(\omega^*)\mu_s(\omega^*) \\
&= \frac{1}{\sum_\omega \mu(\omega) f(s|\omega)} \left(\sum_{\omega<\omega^*} D_{a',a}(\omega)\mu(\omega) f(s|\omega) + D_{a',a}(\omega^*)\mu(\omega^*) f(s|\omega^*) \right) \\
&\geq \frac{\mu(\omega^*) f(s|\omega^*)}{\sum_\omega \mu(\omega) f(s|\omega)} \left(-\overline{D} \frac{\sum_{\omega<\omega^*} \mu(\omega) f(s|\omega)}{\mu(\omega^*) f(s|\omega^*)} + D_{a',a}(\omega^*) \right).
\end{aligned}$$

DUB implies directional distinguishability (Lemma 1), and so there is a signal $\overline{s}$—or more precisely, a set of such signals of positive probability—such that

$$\frac{\sum_{\omega<\omega^*} \mu(\omega) f(\overline{s}|\omega)}{\mu(\omega^*) f(\overline{s}|\omega^*)} < \frac{D_{a',a}(\omega^*)}{\overline{D}},$$

which implies $\mathbb{E}_{\mu_{\overline{s}}}[D_{a',a}(\omega)] > 0$. Since $a \in c(\mu)$ was arbitrary, it follows that μ is not stationary. *Q.E.D.*

Proof of Theorem 1 part 2. The result is subsumed by Proposition 1. *Q.E.D.*

Proof of Theorem 1 part 3. Suppose Ω is finite. By Lemma 1, a failure of DUB implies a failure of directional distinguishability. Assume there is some state ω^* that is not distinguishable from $\{\omega : \omega < \omega^*\}$. (The argument when ω^* is not distinguishable from $\{\omega : \omega > \omega^*\}$ is symmetric.) This means that for some $\varepsilon > 0$ and some belief $\mu \in \Delta\{\omega' : \omega' \leq \omega^*\}$ with $\mu(\omega^*) > 0$, it holds that[26]

$$\forall s : \frac{\sum_{\omega'<\omega^*} \mu(\omega') f(s|\omega')}{\mu(\omega^*) f(s|\omega^*)} > \varepsilon. \tag{2}$$

Take an arbitrary pair of actions a' and a^*. Consider any SCD utility u satisfying $D_{a',a^*}(\omega) \equiv u(a',\omega) - u(a^*,\omega) = 1$ for all $\omega < \omega^*$, and $D_{a',a^*}(\omega^*) = -\varepsilon$. So a' is strictly preferred to a^* in

[26] To be fully precise, the following inequality need only hold for almost all signals when S is an interval; we suppress such qualifiers here and in other proofs.

all $\omega < \omega^*$, while a^* is strictly preferred to a' at ω^*. It follows that

$$\forall s : \mathbb{E}_{\mu_s}[D_{a',a^*}(\omega)] = \frac{\left(\sum_{\omega<\omega^*} \mu(\omega) f(s|\omega)\right) - \varepsilon\mu(\omega^*) f(s|\omega^*)}{\sum_{\omega} f(s|\omega)\mu(\omega)} > 0,$$

where the strict inequality is by (2). That is, if the prior is μ, then no matter the signal, a' is strictly preferred to a^*. The belief μ is thus stationary (either because the choice set is $\{a', a^*\}$, or for larger choice sets because u is such that all other actions are strictly dominated) even though it does not have adequate knowledge. There is thus inadequate learning. *Q.E.D.*

Proof of Proposition 1. Suppose that u fails SCD. Then there is a choice set $A = \{a', a''\}$ and states $\omega_1 < \omega_2 < \omega_3$ such that $\min\{D_{a',a''}(\omega_1), D_{a',a''}(\omega_3)\} > 0 > D_{a',a''}(\omega_2)$, where $D_{a',a''}(\omega) \equiv u(a', \omega) - u(a'', \omega)$. Plainly, any belief whose support includes $\{\omega_1, \omega_2, \omega_3\}$ does not have adequate knowledge.

For any signal structure f, any signal s, and any belief $\mu \in \Delta\{\omega_1, \omega_2, \omega_3\}$,

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

Under MLRP, the strictly positive functions

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

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

$$\mu(\omega_1)h(s) + \mu(\omega_3)g(s) = \frac{1-\varepsilon}{2}\left(h(s) + g(s)\right) \geq \frac{1-\varepsilon}{2}\delta > \varepsilon = \mu(\omega_2). \qquad (4)$$

Combining (4) and the equivalences leading up to (3), we have that for all s, $\mathbb{E}_{\mu_s}[D_{a',a''}(\omega)] > 0$. The belief μ is thus stationary but does not have adequate knowledge, and so there is inadequate learning. *Q.E.D.*

The next result will be used in the proof of Proposition 2. It is a slight strengthening of the standard result that (strict) MLRP implies that posteriors after any two signals are ordered by (strict) first-order stochastic dominance, FOSD hereafter. We say that signals s and s' are *equivalent* if there is a constant c such that for all ω, $f(s|\omega) = cf(s'|\omega)$. That is, equivalent signals generate the same posteriors from any prior.

Lemma 2. *Assume the signal structure f satisfies MLRP. For any full-support prior μ and non-equivalent signals $s < s'$, the posterior $\mu_{s'}$ strictly dominates the posterior μ_s in the sense of FOSD: $\forall \omega < \sup \Omega$, $\sum_{\omega' \leq \omega} \mu_s(\omega') > \sum_{\omega' \leq \omega} \mu_{s'}(\omega')$.*

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

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

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

$$\sum_{\omega' \leq \tilde{\omega}} \mu_s(\omega') = \sum_{\omega' \leq \tilde{\omega}} \mu_{s'}(\omega'). \tag{5}$$

As s and s' are not equivalent, there is some ω such that $\mu_s(\omega) \neq \mu_{s'}(\omega)$.[27] We proceed assuming $\omega \leq \tilde{\omega}$; a symmetric argument applies if $\omega > \tilde{\omega}$. Equation 5 implies that there exists $\omega^* \leq \tilde{\omega}$ such that $\mu_s(\omega^*) < \mu_{s'}(\omega^*)$. For any ω' with $\omega' > \tilde{\omega} \geq \omega^*$,

$$\frac{\mu_s(\omega^*)}{\mu_s(\omega')} = \frac{\mu(\omega^*)f(s|\omega^*)}{\mu(\omega')f(s|\omega')} \geq \frac{\mu(\omega^*)f(s'|\omega^*)}{\mu(\omega')f(s'|\omega')} = \frac{\mu_{s'}(\omega^*)}{\mu_{s'}(\omega')},$$

where the inequality is by MLRP. Hence, for any $\omega' > \tilde{\omega} \geq \omega^*$, $\frac{\mu_{s'}(\omega')}{\mu_s(\omega')} \geq \frac{\mu_{s'}(\omega^*)}{\mu_s(\omega^*)} > 1$, and in particular, $\mu_{s'}(\omega') > \mu_s(\omega')$. So $\sum_{\omega' > \tilde{\omega}} \mu_{s'}(\omega') > \sum_{\omega' > \tilde{\omega}} \mu_s(\omega')$, a contradiction to Equation 5. *Q.E.D.*

Proof of Proposition 2. Take any signal structure f that satisfies MLRP but fails DUB. We begin by constructing an SCD preference u with useful properties, and then argue that given f and u, there is inadequate learning for any full-support prior.

We use the fact that the failure of DUB under MLRP implies a failure of directional distinguishability (by Lemma 1, MLRP is irrelevant for this point when Ω is finite; Corollary 1 in Appendix B establishes that directionality distinguishability also implies DUB when Ω is infinite under MLRP). Let ω^* be a state that is not distinguishable from its lower set

[27] If $\mu_s(\omega) = \mu_{s'}(\omega)$ for all ω, then Bayes rule implies that for all ω, $f(s|\omega) = f(s'|\omega)\frac{\sum_{\omega'} \mu(\omega')f(s|\omega')}{\sum_{\omega'} \mu(\omega')f(s'|\omega')}$; hence, s and s' are equivalent signals.

$\{\omega : \omega < \omega^*\}$; the argument for the other case is similar. So there exists $\mu \in \Delta\{\omega' : \omega' \leq \omega^*\}$ with $\mu(\omega^*) > 0$ and a small $\varepsilon > 0$ such that for every signal s, $\frac{\sum_{\omega<\omega^*} \mu(\omega) f(s|\omega)}{\mu(\omega^*) f(s|\omega^*)} \geq \varepsilon$. By taking the conditional distribution of μ on $\{\omega : \omega < \omega^*\}$, call it μ', and $z \equiv \frac{\varepsilon\mu(\omega^*)}{1-\mu(\omega^*)} \in (0, 1)$, we obtain

$$\forall s : \sum_{\omega<\omega^*} \mu'(\omega) f(s|\omega) \geq z f(s|\omega^*). \tag{6}$$

Consider an SCD preference u such that action a' is strictly preferred in all states below ω^*, while a^* is strictly preferred in ω^* and higher states. In particular, take $D_{a',a^*}(\omega) \equiv u(a', \omega) - u(a^*, \omega) = 1$ for all $\omega < \omega^*$, and $D_{a',a^*}(\omega) = -1$ for all $\omega \geq \omega^*$. Hence, under uncertainty, an agent strictly prefers a' to a if and only if $\Pr(\{\omega : \omega < \omega^*\}) > 1/2$, and is indifferent if and only if this probability is $1/2$.

Fix any full-support prior, the choice set $A = \{a', a^*\}$, and an equilibrium σ. It is without loss of generality to assume that for any history h, $\sigma(a'|s, h)$ is decreasing (weakly) in s. To see that, begin by noting that Lemma 2 implies that for non-equivalent signals $s < s'$, $\sigma(a'|s, h) > \sigma(a'|s', h)$, with either $\sigma(a'|s, h) = 1$ or $\sigma(a'|s', h) = 0$. Thus, the set of signals after which an agent randomizes must all be equivalent to each other. Since the agent is indifferent when randomizing, we can assume a constant randomization probability for all those signals without changing the public belief that is generated from a' being chosen at h.

Now assume, towards contradiction, that there is adequate learning. So there is an infinite history with a herd on a', $h^\infty \equiv (\underbrace{a_1, \ldots, a_{m-1}}_{\text{each } a_i \in \{a', a^*\}}, a', a', \ldots)$ that occurs with positive probability in state $\omega^* - 1$. Let μ^∞ denote the limit public belief resulting from this history. We claim that $\operatorname{Supp} \mu^\infty = \{\omega : \omega < \omega^*\}$. To see why, first note that adequate learning implies that a herd on a' cannot occur in any state at which a' is not optimal. So $\operatorname{Supp} \mu' \subseteq \{\omega : c(\omega) = a'\} = \{\omega : \omega < \omega^*\}$. Given a finite sub-history h^n of h^∞, let $\Pr(a'|h^n, \omega) \equiv \sum_s \sigma(a'|s, h^n) f(s|\omega)$ be the probability that agent n plays action a' when the state is ω. The probability of playing a' is decreasing in ω because $\sigma(a'|s, h^n)$ is decreasing in s and f satisfies MLRP. In particular,

$$\forall \omega < \omega^* : \Pr(a'|h^n, \omega) \geq \Pr(a'|h^n, \omega^* - 1) > 0,$$

and since h^∞ has positive probability at $\omega^* - 1$, it follows that

$$\forall \omega < \omega^* : \prod_{n=m}^{\infty} \Pr(a'|h^n, \omega) \geq \prod_{n=m}^{\infty} \Pr(a'|h^n, \omega^* - 1) > 0. \tag{7}$$

On the other hand, since $\Pr(h^\infty|\omega^*) = 0$ but $\Pr(h^m|\omega^*) > 0$ (as every finite history has positive probability in all states), we have

$$\prod_{n=m}^{\infty} \Pr(a'|h^n, \omega^*) = 0. \tag{8}$$

For the belief $\mu' \in \Delta\{\omega : \omega < \omega^*\}$ and $z \in (0,1)$ both identified in (6), it holds that

$$\begin{aligned}
\sum_{n=m}^{\infty} \log(1 - z\Pr(a^*|h^n, \omega^*)) &\geq \sum_{n=m}^{\infty} \log\left(1 - \sum_{\omega<\omega^*} \mu'(\omega)\Pr(a^*|h^n,\omega)\right) \qquad \text{(using (6))}\\
&= \sum_{n=m}^{\infty} \log\left(\sum_{\omega<\omega^*} \mu'(\omega)\Pr(a'|h^n,\omega)\right)\\
&\geq \sum_{n=m}^{\infty}\sum_{\omega<\omega^*} \mu'(\omega)\log(\Pr(a'|h^n,\omega)) \qquad \text{(by Jensen's inequality)}\\
&= \sum_{\omega<\omega^*} \mu'(\omega)\sum_{n=m}^{\infty} \log(\Pr(a'|h^n,\omega)) \qquad \text{(by Tonelli's theorem)}\\
&= \sum_{\omega<\omega^*} \mu'(\omega)\log\left(\prod_{n=m}^{\infty} \Pr(a'|h^n,\omega)\right)\\
&> -\infty \qquad \text{(by (7))}.
\end{aligned} \tag{9}$$

Below we will invoke the mathematical fact that for arbitrary sequences (S_n) and (T_n) and constant $c > 0$, if $\lim_{n\to\infty} \frac{S_n}{T_n} = c > 0$ and $\sum_n S_n < \infty$, then $\sum_n T_n < \infty$.[28] Let $S_n = -\log(1 - z\Pr(a^*|h^n,\omega^*))$ and $T_n = -\log(1-\Pr(a^*|h^n,\omega^*))$. 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 (9) implies $\lim_{n\to\infty}\Pr(a^*|h^n,\omega^*) = 0$. The aforementioned mathematical fact implies that

$$\sum_{n=m}^{\infty} -\log(1-\Pr(a^*|h^n,\omega^*)) < \infty.$$

As $\Pr(a'|h^n,\omega^*) = 1 - \Pr(a^*|h^n,\omega^*)$, it further follows that

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

[28] 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 $\sum_n T_n \leq \sum_{n\leq N} T_n + \sum_{n>N} S_n/c' < \infty$.

which contradicts (8). *Q.E.D.*

Proof of Proposition 3. We provide an argument to establish part 1 of DUB. A symmetric argument can be used to establish part 2. We will use the following claim.

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

Proof of Claim: Suppose not, to contradiction. Then for some $\overline{s} > 0$ and $\varepsilon \in (0,1)$, every $s \geq \overline{s}+1$ has $k_s > 0$ such that either (i) $k_s \geq 1/\overline{s}$ and $\frac{g(s+k_s)}{g(s)} > \varepsilon$, or (ii) $0 < k_s < 1/\overline{s}$ and $\frac{g(s+k_s)}{g(s)} > 2$. Hence, for all $s \geq \overline{s}+1$, either (i) $\frac{g(s+k_s)}{g(s)} > \varepsilon \geq \varepsilon^{k_s\overline{s}}$ (because $k_s\overline{s} \geq 1$), or (ii) $\frac{g(s+k_s)}{g(s)} > 2 > \varepsilon^{k_s\overline{s}}$ (because $\varepsilon < 1$). Moreover, there is an increasing sequence $(s_i)_{i=1}^{\infty}$ such that $s_1 = \overline{s}+1$ and for all $i > 1$, $s_i = s_{i-1} + k_{s_{i-1}}$. Note that for all i, $s_i = (\overline{s}+1) + \sum_{j=1}^{i-1} k_{s_j}$.

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

$$(s_i - \overline{s} - 1)\overline{s}\log(\varepsilon) + \log(g(\overline{s})) \leq \log(g(s_i)). \tag{10}$$

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

$$\log(g(s_i)) \leq -(s_i)^p. \tag{11}$$

The left-hand side of inequality (10) is linear in s_i while the right-hand side of inequality (11) has exponent $p > 1$, so for large enough s_i these inequalities are in contradiction.

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

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

Finally, to verify DUB part 1, fix any state ω and define $\overline{s}_i \equiv s_i^* + \omega$. Then, for all i and $\omega' < \omega$,

$$\frac{f(\overline{s}_i|\omega')}{f(\overline{s}_i|\omega)} = \frac{g(\overline{s}_i - \omega')}{g(\overline{s}_i - \omega)} = \frac{g(s_i^* + (\omega - \omega'))}{g(s_i^*)} \leq 2,$$

and, since $s_i^* \to \infty$, for any sufficiently large i,

$$\frac{f(\overline{s}_i|\omega')}{f(\overline{s}_i|\omega)} = \frac{g(s_i^* + (\omega - \omega'))}{g(s_i^*)} \leq \sup_{k \geq 1/s_{i-1}^*} \frac{g(s_i^* + k)}{g(s_i^*)} \leq \frac{1}{i-1} \to 0,$$

establishing part 1 of DUB's requirement. *Q.E.D.*

B. DUB vs. Directional Distinguishability when $|\Omega| = \infty$

Although DUB characterizes directional distinguishability when Ω is finite, the latter is more permissive when Ω is infinite (cf. Lemma 1). This appendix elaborates on the gap between these two definitions; we also establish that the gap ceases to exist under MLRP.

The gap between DUB and directional distinguishability has to do with the uniform boundedness required by DUB, i.e., point (ii) in each part of Definition 2. We first explain why this requirement is there. Given a state ω and a prior μ with $\mu(\omega) > 0$, take a sequence of signals $(s_i)_{i=1}^{\infty}$ witnessing that ω is distinguishable from its lower set $\{\omega' : \omega' < \omega\}$:

$$\lim_{i \to \infty} \frac{\sum_{\omega' < \omega} \mu(\omega') f(s_i|\omega')}{\mu(\omega) f(s_i|\omega)} = 0.$$

Clearly, $\frac{f(s_i|\omega')}{f(s_i|\omega)} \to 0$ pointwise for every $\omega' < \omega$, i.e., the sequence simultaneously verifies that ω is (pairwise) distinguishable from each ω'. So point (i) in DUB is necessary for directional distinguishability. However, the converse can fail: such pointwise convergence does not guarantee that ω is distinguishable from its lower set. Here is an example.

Example 2. Let $\Omega = S = \mathbb{Z}$ and take the probability mass function $g(x) = ke^{-x^2}$, with $k > 0$ a normalizing constant. Let $\lambda \in (0, 1)$ and define the probability mass function

$$f(s|\omega) = \begin{cases} (1-\lambda)g(s-\omega) & \text{if } s \neq -\omega \\ (1-\lambda)g(s-\omega) + \lambda & \text{if } s = -\omega. \end{cases}$$

This signal structure has two notable properties. First, for any $\omega' < \omega$,

$$\lim_{s \to \infty} \frac{f(s|\omega')}{f(s|\omega)} = \lim_{s \to \infty} \frac{g(s-\omega')}{g(s-\omega)} = 0,$$

Second, for any ω,

$$\lim_{s \to \infty} \frac{f(s|-s)}{f(s|\omega)} = \lim_{s \to \infty} \left[\frac{g(2s)}{g(s-\omega)} + \frac{\lambda}{(1-\lambda)g(s-\omega)} \right] = \infty.$$

The first point, which verifies the pointwise convergence to 0 discussed before the example, says that given any two states, sufficiently large signals provides near certainty about the higher state relative to the lower one. The second point says that at the same time, given any state ω, a large signal s also provides near certainty about state $-s$ relative to ω.

Now consider the prior μ given by

$$\mu(\omega) = \begin{cases} k'e^{\omega} & \text{if } \omega \leq 0 \\ 0 & \text{if } \omega > 0, \end{cases}$$

where $k' > 0$ is a normalizing constant. Denoting the posterior after signal s by μ_s, we compute

$$\lim_{s\to\infty} \frac{\sum_{\omega<0} \mu(\omega) f(s|\omega)}{\mu(0) f(s|0)} \geq \lim_{s\to\infty} \frac{\mu(-s) f(s|-s)}{\mu(0) f(s|0)} \geq \lim_{s\to\infty} \frac{\mu(-s)}{\mu(0)} \frac{\lambda}{g(s)} = \lim_{s\to\infty} \frac{\lambda e^{s^2-s}}{k} = \infty.$$

That is, large signals make the agent near certain the state is not 0 — even though large signals provide near certainty that the state is 0 relative to any given state $\omega' < 0$. Consequently, directional distinguishability fails: state 0 is not distinguishable from its lower set.[29] $\square$

Directional distinguishability fails in Example 2 because given any state ω the ratio $f(s|\omega')/f(s|\omega)$ is unbounded over s and ω', even though the ratio goes to 0 when either s or ω' is fixed and only the other goes to ∞ or $-\infty$, respectively. It is to avoid such issues that DUB imposes the uniform boundedness requirement in point (ii) of each part of Definition 2.

One may ask whether the uniform boundedness requirement can be weakened while still guaranteeing directional distinguishability, and if so, how. We do not have a precise answer. However, we can establish that the uniform boundedness is necessary for the following rather natural strengthening of directional distinguishability.

Definition 5. There is *prior-independent directional distinguishability* (PIDD), if for every ω there are two sequences $(\overline{s}_i)_{i=1}^{\infty}$ and $(\underline{s}_i)_{i=1}^{\infty}$ such that for every $\mu \in \Delta\Omega$ with $\mu(\omega) > 0$,

$$\lim_{i\to\infty} \frac{\sum_{\omega'<\omega} \mu(\omega') f(\overline{s}_i|\omega')}{\mu(\omega) f(\overline{s}_i|\omega)} = 0 \quad \text{and} \quad \lim_{i\to\infty} \frac{\sum_{\omega'>\omega} \mu(\omega') f(\underline{s}_i|\omega')}{\mu(\omega) f(\underline{s}_i|\omega)} = 0. \tag{12}$$

[29] Strictly speaking, we have only argued that no sequence of signals going to ∞ can distinguish state 0 from its lower set. But note that for any $\omega < 0$, $\frac{f(s|\omega)}{f(s|0)} = \frac{g(s-\omega)}{g(s)}$ is decreasing in s over the set $S \setminus \{-\omega\}$, so no sequence of signals that is bounded above will work either.

PIDD requires that for each ω there is a single sequence of signals that provides certainty about ω relative to its lower states regardless of the prior, and similarly for higher states. We will shortly see that such prior independence arises naturally (Corollary 1). Moreover, verifying directional distinguishability absent PIDD would be daunting. In any case, the following result establishes that DUB characterizes PIDD.

Proposition 5. *The signal structure has DUB if and only if it has PIDD.*

Proof. In the proof of Lemma 1, the argument that DUB $\implies$ directional distinguishability used a signal sequence $(\bar{s}_i)_{i=1}^{\infty}$ that depends only on the state ω but not the prior $\mu \in \Delta\Omega$. Hence, that argument actually establishes DUB $\implies$ PIDD.

We now prove that PIDD $\implies$ DUB. Suppose PIDD holds, and to contradiction, that DUB does not hold. For concreteness, suppose there is ω such that the signal sequence $(\bar{s}_i)_{i=1}^{\infty}$ witnessing PIDD of ω from its lower states must not satisfy part 1 of DUB's definition. (The argument is analogous if the sequence if it is part 2 that fails with the sequence $(\underline{s}_i)_{i=1}^{\infty}$ from PIDD.)

By considering any full-support prior μ in PIDD's property (12), we see that for all $\omega' < \omega$, $\lim_{i\to\infty} \frac{f(\bar{s}_i|\omega')}{f(\bar{s}_i|\omega)} = 0$. So it must be that the uniform boundedness in DUB that fails. This means that for any integer $k \geq 1$ we can find an index n_k and a state $\omega_k < \omega$ such that $\frac{f(\bar{s}_{n_k}|\omega_k)}{f(\bar{s}_{n_k}|\omega)} > 2^{k+1}$, and moreover, $n_k \to \infty$ as $k \to \infty$. For any prior μ such that $\mu(\omega) > 0$ and, for all $k \geq 1$, $\mu(\omega_k) = 1/2^{k+1}$, it follows that for all $k \geq 1$,

$$\frac{\sum_{\omega'<\omega} \mu(\omega') f(\bar{s}_{n_k}|\omega')}{\mu(\omega) f(\bar{s}_{n_k}|\omega)} > \frac{\mu(\omega_k) f(\bar{s}_{n_k}|\omega_k)}{\mu(\omega) f(\bar{s}_{n_k}|\omega)} > \frac{1}{\mu(\omega)}.$$

But according to PIDD, the left-most fraction in the display becomes arbitrarily small when n_k is large enough, a contradiction. *Q.E.D.*

Finally, we observe that under MLRP, no matter ω or the prior, only a sequence $(\bar{s}_i)_{i=1}^{\infty}$ that goes to $\sup S$ and a sequence $(\underline{s}_i)_{i=1}^{\infty}$ that goes to $\inf S$ can satisfy (12). It follows that:

Corollary 1. *If the signal structure has MLRP, then PIDD (and hence DUB) is equivalent to directional distinguishability.*

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," *The 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 (2019): "A General Analysis of Sequential Social Learning," Forthcoming in *Mathematics of Operations Research*.

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.

BANERJEE, A. V. (1992): "A Simple Model of Herd Behavior," *The 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.

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.

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

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

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

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

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 (2019): "Single-Crossing Differences on Distributions," Unpublished.

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

LOBEL, I. AND E. SADLER (2015): "Information diffusion in networks through social learning," *Theoretical Economics*, 10, 807–851.

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

MILGROM, P. R. (1981): "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12, 380–391.

PARK, A. AND H. SABOURIAN (2011): "Herding and Contrarian Behavior in Financial Markets," *Econometrica*, 79, 973–1026.

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.

SMITH, L., P. N. SØRENSEN, AND J. TIAN (2021): "Informational Herding, Optimal Experimentation, and Contrarianism," Forthcoming in *Review of Economic Studies*.