Information Cascades and Social Learning

Sushil Bikhchandani
David Hirshleifer
Omer Tamuz
Ivo Welch

August 18, 2022

Abstract

In an *information cascade*, an agent who observes others chooses the same action regardless of her own private information signal. Cascades result in poor information aggregation, inaccurate decisions, and fragility of mass behaviors. We review the theory of information cascades and social learning, and discuss important themes, insights and applications of this literature as it has developed over the last thirty years. We also highlight open questions and promising directions for further theoretical and empirical exploration.

We thank the editor, Steven Durlauf, and four anonymous referees for very helpful comments.

Keywords: Information Cascades, Social Learning, Herding, Fads, Fashion, Conformity, Culture

JEL Codes: D8 Information, Knowledge, and Uncertainty, D7 Analysis of Collective Decision-Making, D9 Micro-Based Behavioral Economics

1 Introduction

People rely heavily on the information of others in forming their opinions and selecting actions. An evolutionary explanation for this is that gleaning useful information from others increased fitness. The updating of beliefs based on observation of the actions of others, observation of the consequences of these actions, or conversation with others, is called *social learning*.

An empirical literature has documented the importance of social learning and its consequences for social outcomes. A basic implication of social learning is that people who observe others' choices often start behaving similarly.

Some examples illustrating this point give a hint of the scope of social learning effects. In an early experiment, confederates of the experimenter stared up at absolutely nothing, and drew crowds of observers doing likewise (Milgram, Bickman and Berkowitz (1969)). In a domain with higher stakes, there is extensive evidence of social influence in the mating decisions of humans and other animals. In one experiment, seeing a rival show interest in a member of the opposite sex caused human observers to rate the individual as more appealing (see Bowers et al. (2012) and the related evidence of Little et al. (2008)). In another, female guppies switched their choice of mate to match the choices made by other females (Dugatkin and Godin (1992)). A life-and-death example is the decision of patients to refuse good kidneys for transplant owing to refusal by previous patients earlier in the queue (Zhang (2010)).

Social learning also helps shape economic activity. For example, in labor markets, unemployment leads to a form of stigma, as employers view gaps in a resume as indicating rejection by other employers (see field experiments by Oberholzer-Gee (2008) and Kroft, Lange and Notowidigdo (2013)). Financial professionals from the Chicago Board of Trade also engage in social learning. For example, in a laboratory experimental setting, professionals from the Chicago Board of trade tend to follow the incorrect actions of others instead of their own signals (Alevy, Haigh and List (2007)).

In the realm of politics, early primary election wins cause later voters to support the winner. For instance, primary election victories of John Kerry over Howard Dean in 2004 generated political momentum and later wins for Kerry in the U.S. presidential primary contest (Knight and Schiff (2010)).

These examples illustrate two patterns. The first is a tendency for agents to converge upon the same action—perhaps even a suboptimal one. Behavioral convergence seems to arise spontaneously, even when there is no opportunity or incentive to punish deviants, and sometimes despite the opposing force of negative payoff externalities. The second is that outcomes are often path-dependent, in the sense of being sensitive to the sequence of information arrival. Examples of path-dependence are the effects of spells of unemployment, of strings of kidney rejections, and of political momentum. Furthermore, outcomes are often fragile in the sense of being sensitive to shocks. Examples of fragility are the rise and fall of surgical fads and quack medical treatments on the part of physicians who rely on their colleagues' practices (Taylor (1979), Robin (1984)).

The elemental fact of social learning raises several important questions. One set of questions concerns efficiency. When people make decisions in sequence and have the opportunity to observe each other's actions or payoffs, or communicate with each other, do they eventually make correct choices? In other words, is dispersed information aggregated effectively? If agents do eventually learn which action is optimal, how quickly does this occur? And does rationality promote better social outcomes?

Another fundamental set of questions concerns the homogeneity and stability of outcomes. Will individuals' actions and their beliefs about the state of the world eventually be the same? If agents do converge upon the same action, how stable is this outcome with respect to exogenous shocks? Is the system prone to sudden shifts, or fads?

There are also important questions about what determines outcomes. For example, how do the costs and distribution of private information signals affect learning and behavior? And what are the effects of the network structure of social observation?

Early papers sought to address these questions by means of the concepts of information cascades and herding (Banerjee (1992), Bikhchandani, Hirshleifer and Welch (1992), and Welch (1992)).[1] Information cascades were soon applied in such fields as anthropology, computer science, economics, law, political science, psychology, sociology, and zoology. Over time, a theoretical literature examined the robustness of the conclusions of the basic information cascades setting to varying different model assumptions. There has also been a flowering of social learning modeling in which the assumptions used are more distant from the basic cascades settings. More recently, an empirical literature (including experimental research) has tested both rational and imperfectly rational social learning models.

In the models of Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992), and Welch (1992) agents observe private information signals and make decisions in sequence based upon observation of the actions of predecessors. A key implication of these models is that, under appropriate conditions, there will always be an *information cascade*: a situation in which an agent or a sequence of agents act independently of their private information signals. This happens when the information derived from social observation overwhelms the given agents' signals. When in a cascade, an agent's action is uninformative to later agents, so that social learning is blocked — at least for a time. As a result, if agents are ex ante identical, then all later agents face the same decision problem and make the same choice.

Cascading imposes an adverse information externality upon subsequent agents. This leads to a surprising result. Even though agents are fully rational and collectively possess very strong information, their decisions are often incorrect. This contrasts with the standard conclusion that reasonably good aggregation of the signals of a large population of informed agents leads to accurate choices. The problem of information

[1] We define information cascades later. Banerjee's term, "herding," has essentially the same meaning as information cascades in Bikhchandani, Hirshleifer and Welch (1992), but "herding" has several different meanings in economics, and even within the social learning literature. We therefore use the term "information cascades" for the concept introduced by these two papers.

externalities is also present in social learning models in which cascades do not occur.

Cascades in the basic setting above are precarious; agents in a cascade are somewhat close to indifferent between two action alternatives. So in these models there is a systematic tendency to reach a resting point at which behavior is sensitive to small shocks. In other words, social outcomes are fragile; the arrival of a small amount of information, or even the possibility of such an arrival, can dislodge agents from a long-standing cascade. So the cascades model offers a possible explanation for why social behaviors are often idiosyncratic and volatile. This is the case even though all individuals are fully rational, and enough information to make the correct decision is available to society as a whole.

Models of social learning, including cascades models, can help explain volatile aggregate behaviors and dysfunctional social outcomes in a range of social and economic domains. In this survey we provide an overview of this research. Our main focus is on agents who engage in Bayesian or quasi-Bayesian updating, and are trying to make good choices (i.e., to optimize).

In § 2 we present what we call the *Simple Binary Model*. This is the simplest model of information cascades, which illustrates some basic intuition behind this phenomenon. It also helps isolate the effects of generalizing the assumptions in various ways. In § 3 and § 4 we explore the robustness of its conclusions to varying assumptions about the action space and the signal distribution, respectively.

In later sections we consider a number of more fundamental deviations from the basic cascades setting. For example, allowing agents to choose whether to act immediately or to delay offers new insight about boom and bust dynamics in investment and market entry contexts (see § 5). In § 6 we relax the assumption that agents observe all of their predecessors and that agents do not observe predecessors' payoffs.

The availability of social information can reduce the incentive to acquire private information. In a sequential social learning setting, fixed costs of information acquisition encourages cascading upon the actions of others, because agents can cheaply follow predecessors, thereby avoiding the cost of acquiring private signals. This in turn reduces the information content of actions for later observers. This is explored in § 7.

In § 8 we consider how payoff externalities between the actions of different agents can either reinforce or oppose the disposition to follow the behavior of predecessors. One important form of payoff interaction occurs when agents participate in market exchange. § 9 considers how endogenous price determination affects the social learning process.

In the basic cascades models, preferences are common knowledge. When people do not know the preferences of others, it becomes harder to deduce predecessors' signals from their actions. The effects of heterogeneous preferences on social learning are considered in § 10.

People tend to observe and learn from those that they are socially connected to. This raises the question of how social outcomes and welfare are shaped by the network of observation and communication. For example, does the information of some subset of individuals end up unduly influencing outcomes for society as a whole? To address such questions, we review research on social learning in networks in § 11. We consider settings in which people have the opportunity to act and learn from each other either once or repeatedly. We consider the role of the geometric structure of the network in determining the social outcome. One key insight, for example, is that networks that are (in some sense) egalitarian are conducive to better social learning outcomes.

One of the key directions covered here is the study of how limited rationality affects social learning (see especially § 6.2, § 11.3.2 and § 10). Allowing for limited rationality is especially important in realistic social networks, in which inference problems are so complex that perfect rationality is implausible. Our coverage of limited rationality focuses on quasi-Bayesian decision makers rather than mechanistic agents.

Finally, in § 12 we cover models of information cascades in a variety of applied domains, including politics, law, product markets, financial markets, and organizational structure. The social learning approach, by focusing on the role of information externalities, offers new perspectives about a wide range of human behaviors.

The analysis of social learning uncovers effects that differ in some ways from most of traditional information economics. Where much of information economics is strategic, in that the actions of some agents directly affect the payoffs of others, in much of the social learning literature, externalities are purely informational, with no direct payoff interactions. Thus, forces of central focus for much of information economics, such as adverse selection and moral hazard, are not the main drivers of most social learning models.

The formal modelling of general social influence began primarily in social sciences other than economics. Interactions via social networks were studied heavily in sociology, as with the influential DeGroot model (DeGroot (1974)), and cultural evolution was studied in anthropology and other fields (Boyd and Richerson (1985)). However, these models typically used mechanistic assumptions about how agents update their actions, traits, or 'opinions,' rather than modeling Bayesian or quasi-Bayesian updating of beliefs and resulting optimal actions. In contrast, the economic approach to social learning allows for agents who are intelligent in their belief updating and process of optimization.

In the 1990s, economists started to focus attention upon social learning in general, rather than just as a by-product of market interactions or game-theoretic interactions between small numbers of agents.[2] A newer element in this literature is the possibility that people observe or talk to others even when the target of observation has no incentive or intent to influence others. Also integral to this literature is that many agents act and observe *sequentially*, so that the learning process iterates, as is often the case in practice.

[2]The latter includes models of learning from market price in financial markets (Grossman and Stiglitz (1976)) and other markets (Akerlof (1976), Spence (1978)). The former includes the "agreeing to disagree" literature (Aumann (1976), Geanakoplos and Polemarchakis (1982)) and the cheap talk literature (Crawford and Sobel (1982)).

Even without learning (the focus of this survey), social interaction can cause behaviors to converge, owing, for example, to payoff externalities or utility interactions (Arthur (1989)), reputation effects (Scharfstein and Stein (1990), Ottaviani and Sørensen (2000)) or a preference for conformity (Bernheim (1994)). Some models also study how social interaction induces dynamics of convergent or divergent behaviors (Kirman (1993)). Our focus is on behavioral convergence or divergence that derives from social learning and information cascades. However, most actual applications involve the interaction of several possible factors, including information, rewards and punishments, and payoff externalities. The integration of social learning and cascades with other considerations has led to a richer palette of theory about the process by which society chooses technologies, ideas, governments, organizational choices, conventions, legal precedents, and market outcomes.

The interested reader may also consult other surveys of social learning (e.g., Gale (1996), Bikhchandani, Hirshleifer and Welch (1998), Chamley (2004*b*), Vives (2010), Acemoglu and Ozdaglar (2011), and Golub and Sadler (2016)). Our review includes coverage of more recent research and an array of topics, including social networks, repeated moves, and psychological bias, while highlighting the role of information cascades in this rapidly evolving field. Complementing our theoretical focus are several reviews with a primarily empirical focus, such as Hirshleifer and Teoh (2009), Anderson and Holt (2008), Sacerdote (2014), and Blume et al. (2011). These surveys cover tests of cascades theory in the experimental laboratory (Anderson and Holt (2008)), in field experiments (Duan, Gu and Whinston (2009)), and with archival data (Tucker, Zhang and Zhu (2013), Amihud, Hauser and Kirsh (2003)).

2 The Simple Binary Model: A Motivating Example

We illustrate several key intuitions in a setting with binary actions, states and signals, which we call the *Simple Binary Model*, hereafter, SBM. The SBM is a modified version of the binary example of Bikhchandani, Hirshleifer and Welch (1992), hereafter BHW.[3] The SBM illustrates how information cascades can block social learning, and can be extended to illustrate many further concepts. Nevertheless, some of the conclusions of the SBM are more robust than others to changes in the model. We recurrently discuss in this survey when these effects do or do not arise in various other social learning settings.

2.1 Basic Setup: Binary Actions, Signals, and States

Individuals $I_1, I_2, I_3, \ldots$ make choices in sequence. Each agent I_n chooses one of two actions, High ($a_n = H$) or Low ($a_n = L$). The underlying state θ, which is not observed, takes one of two possible values, H or L. The two states are equally likely ex ante. We will often view the H state as one in which there is high payoff to some activity and the L state in which there is low payoff to that activity. In such cases we call action H "Adopt," and action L "Reject." Table 1 summarizes the notation used in this survey.

Each agent I_n receives a binary symmetric private information signal $s_n = h$ or $s_n = \ell$, with probability $p := \mathbb{P}[s_n = h|\theta = H] > 1/2$ and $\mathbb{P}[s_n = h|\theta = L] = 1 - p$. Signals are independent conditional on θ.

We refer to p as the *signal precision*. Let the *belief precision* of a possible belief q about the state θ be $|q - 0.5|$. The greater the signal precision, the closer the belief is to 0 or 1, implying more certainty about the state.

By Bayes' rule, the posterior probability $\mathbb{P}[\theta = H|s_n = h] = p$, and $\mathbb{P}[\theta = H|s_n = \ell] = 1 - p$. By the symmetry of the SBM, h and ℓ signals have offsetting effects on posterior beliefs, so $\mathbb{P}[\theta = H|s_1 = h, s_2 = \ell] = 1/2$, and if an agent sees or infers certain numbers of predecessor h and ℓ signals, the updated belief depends only on the difference between these numbers.

[3] The simple binary model is the special case of the BHW model with symmetric binary signals. The model of Banerjee (1992) differs in several substantive ways, as described in § 3.

Table 1: Notation Guide

	Notation	Base case assumption
State	θ	$\theta \in \{L, H\}$
Signal	s_n for agent I_n	$s_n \in \{\ell, h\}$
Action	a_n for agent I_n	$a_n \in \{L, H\}$
Utility	$u(\theta, a)$	$u(\theta, a) = \begin{cases} 1 & \text{if } \theta = a \\ 0 & \text{otherwise} \end{cases}$

All agents have the same utility function $u(\theta, a)$, which is equal to 1 if $a = \theta$ and to 0 otherwise. Therefore, each agent chooses the action that is more likely to match the state given her information. If I_n assigns equal probabilities to both states, she is indifferent between the two actions, in which case we assume that she follows her private signal, $a_n = s_n$.[4]

We refer to *social information* in this survey as information derived from observing others, the *social belief* at any step in the sequence as the belief that an outside observer would have based only on the social information of agent I_n, and an agent's *private belief* as the belief implied solely by the agent's private signal. Each agent's action choice is of course based on the agent's full information set. It is common knowledge that each I_n knows the decision model and information structure of predecessors.

[4]The qualitative properties of this model are robust to changes in the tie-breaking rule. For the purpose of providing minimal examples, we sometimes employ different tie-breaking conventions for the behavior of indifferent agents. Although pedagogically convenient, such ties could be eliminated, with similar results, by slightly perturbations of the model parameters.

We contrast two regimes for social observation:

1. **The Observable Signals Regime:** Each agent observes both the signals and the actions of all predecessors.[5]

2. **The Only-Actions-Observable Regime:** Each agents observes the actions but not the private signals of all predecessors.

In the Observable Signals Regime, the pool of social information always expands, and by the law of large numbers the social belief becomes arbitrarily close to certainty about the correct action—Adopt if $\theta = H$, Reject if $\theta = L$—and thus eventually behave alike. An outcome in which all agents behave alike forever is called a *herd*. We define herds more formally in § 4.3.5.

In the Only-Actions-Observable Regime, each agent's action depends on the publicly observable action history and the agent's own private signal. In this regime, the precision of the social belief is weakly increasing with n, so it is tempting to conjecture that highly accurate outcomes will again be achieved. But in fact, the precision of the social belief hits a finite ceiling, as we will discuss. It is easy to show that the choices of a few early predecessors determine the actions of all later agents. In consequence, agents still herd — but often upon the wrong action, the choice that yields a lower payoff.

Throughout this review, our default premise will be the Only-Actions-Observable regime; we only mention assumptions when they deviate from the base set of assumptions of the SBM. The SBM, and modest variations thereof, are rich sources of insight into social learning.

[5] Actions do not convey any additional information as a predecessor's signal is a sufficient statistic for her action.

2.2 Why information stops accumulating: Information cascades

To see why outcomes are inefficient in the Only-Actions-Observable Regime, consider each agent in sequence. Figure 1 shows the possible chains of events. Agent I_1, Ann, adopts if her private signal is h and rejects if it is ℓ. Agent I_2, Bob, and all successors can infer Ann's signal perfectly from her decision. So if Ann adopts, and Bob's private signal is h, he also adopts. Bob knows that there were two h signals: he infers one from Ann's actions and has observed one privately. If Bob's signal is ℓ, it exactly offsets Ann's h, making him indifferent between adopting and rejecting. By our tie-breaking convention, Bob follows his own signal and rejects. In this setting, regardless of which signals Ann and Bob receive, each chooses an action that matches his or her signal, and thus their actions reveal their signals to later agents.

Agent I_3, Carol, now faces one of three possible situations: (1) Ann and Bob both adopted (their actions were HH), (2) both rejected (LL), or (3) one adopted and the other rejected (HL or HL). In the HH case, Carol adopts regardless of which signal she received, since the majority of observed or inferred signals is h even if she received an ℓ. In other words, Carol's action is independent of her private information signal—she is in an information cascade, defined as follows:

Definition 1. *An agent is said to be in an* information cascade *if it is optimal for her to choose an action which is independent of her private signal.*

Crucially, in the HH case, Carol's action provides no information about her signal to her successors. Her action has not improved the public pool of information; the social belief remains unchanged. Hence all later agents face the same decision problem that she did. Once Carol is in an Adopt cascade, all her successors also adopt, ultimately based only on the observed actions of Ann and Bob. Similarly, in the second case, two Rejects (LL) put Carol and all subsequent agents into a Reject cascade.

In the third case, Ann has adopted and Bob has rejected, or vice versa. Carol infers that Ann and Bob observed opposite signals, so the social belief is that the two

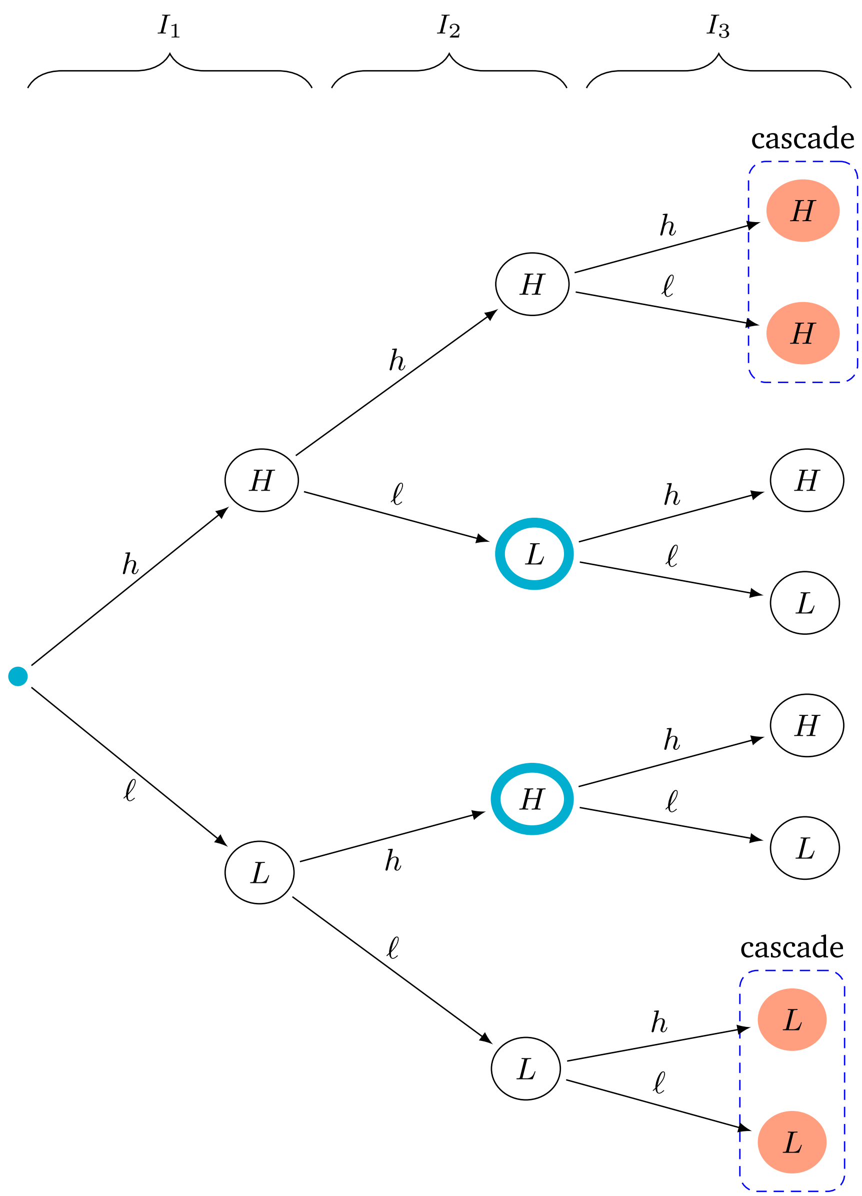

Figure 1: Event Sequence

This figure shows a sequence of possible early events and optimal choices in the SBM. Possible signals are h and ℓ, which are indicative of states H and L respectively. The main focus is on events after an initial h signal. Multiple nodes that are in the same information set from the viewpoint of later observers are surrounded by dashed lines. Action H is correct in state H, and action L is correct in state L. Two circles for I_2 are in cyan to indicate that when these nodes are reached, the next player, I_3, will face the same inference problem as I_1 following the initial cyan node. Nodes that are in an information cascade are shaded in peach.

states are equally likely. Her decision problem is therefore identical to that of Ann, so Carol's decision is to follow her private signal. In turn, this makes the decision problem of agent I_4, Dan, isomorphic to Bob's—he need only pay attention to his latest predecessor.

More generally, taking agents pairwise starting from (I_1, I_2), the only way to avoid an information cascade is for each pair to contain one Adopt and one Reject. If this occurs through any even number of agents, the next agent can infer that the number of past h and ℓ signals was equal, so that past actions can be ignored. Following any such history of paired opposing actions, the next two agents may both receive the same signal (which occurs with probability $p^2 + (1-p)^2$). When the next two agents observe the same signals, they take the same action, so the next agent is in a cascade. Once an agent is in a cascade, so are all successors.[6] It follows that a herd also occurs. Furthermore, since each pair has a fresh chance to start a cascade, a cascade happens eventually, almost surely.

Cascades tend to start early—based on a small preponderance of evidence in the social belief. Even in the least cascade-favorable scenario of very noisy signals ($p \approx 1/2$), by the time the 20$^{\text{th}}$ agent has to decide, the probability of not being in a cascade is already under 0.1%, and the expected number of agents who act on private information before a cascade starts is under four. So the private information of all but a few agents is lost to the group.

Once a cascade starts, no new information about state becomes public; the accuracy of the social belief plateaus. An early preponderance of either Adopts or Rejects causes all subsequent agents to rationally ignore their private signals. These signals thus never join the public pool of knowledge. Agents disregard their private signals before the social belief becomes very accurate. Specifically, as soon as the social belief is more precise than the signal of just a single agent, the next agent falls into a cascade. We will call this *the logic of information cascades*.[7] The logic of information cascades

[6]In general, however, a cascade can occur for just one agent or can continue for any given number of agents.

[7]In a setting with more than two possible signal values, two identical successive actions do not

implies that cascades are often incorrect—they are *idiosyncratic*.

The accuracy improvement from social observation in the Only-Actions-Observable Regime falls far short of the outcome in the Observable Signals Regime, in which eventually everyone makes the correct decision. Indeed, when private signals are very noisy (e.g., $p = 1/2 + \epsilon$), the social outcome is almost pure noise. This falls far short of perfect information aggregation, wherein decisions for later agents would become almost perfectly accurate. Specifically, the increase in accuracy that agents obtain from being able to observe the actions of predecessors is negligible.[8]

The social outcome in the Only-Actions-Observable Regime depends heavily on the order in which signals arrive. If signals arrive in the order $hh\ell\ell\ldots$, then a cascade starts and everyone adopts. If, instead, the same set of signals arrives in the order $\ell\ell hh\ldots$, everyone rejects. So cascades and welfare are *path dependent*.

2.3 Lessons of the Binary Model

The SBM illustrates a number of key concepts that are empirically testable, and which also obtain in some more general settings.

Conformity: People end up following the behavior of others.

> The conclusion that agents conform to the actions of others—even incorrect actions—holds in a wide array of social learning models. The SBM provides a useful benchmark for assessing which assumptions are crucial for this conclusion. In later sections we will see that a similar conclusion can hold with more general

necessarily start a cascade. The general point is that a fairly low-precision social belief can be enough to trigger a cascade.

[8] In the Only-Actions-Observable Regime, the probability of a correct cascade (Adopt if and only if $\theta = H$) can be shown to be $\frac{p^2}{p^2+(1-p)^2}$. For a noisy signal as above, this is approximately $1/2 + 2\epsilon$. This is not much better than an agent could do based solely on private information, which results in a probability of choosing correctly of $p = 1/2 + \epsilon$.

action spaces and signal distributions, and under different assumptions about the timing of moves, observability of others, and the decision of whether to acquire information. However, some variations of modeling assumptions such as unbounded private signals, certain psychological biases, and negative payoff externalities (such as congestion effects) can break this conclusion, even when information cascades still occur.

Idiosyncrasy and Path Dependence: Despite a wealth of private information, which in the aggregate would assure the correct action if it could be made fully available for use, the behavior of most agents often ends up being incorrect—idiosyncrasy. Later in this survey we examine the robustness of this conclusion. Specifically, when does society converge to correct behavior? And under what conditions is there persistent idiosyncrasy?

In particular, the actions of a few early agents tend to be decisive in determining the actions and success of a large numbers of successors. Social outcomes have low predictability. In other words, outcomes are path-dependent. This problem raises the question of whether policy can improve outcomes. As such, there has been extensive interest in path dependence even in settings without social learning (Arthur (1989), Liebowitz and Margolis (1995)).

Information Externality: Each agent takes an action that is individually optimal, without consideration of the informational benefits to later agents. A greater benefit could be conferred upon later agents if early agents were instead to act in ways that reveal their own signals. The wastage of private information blocks efficient outcomes.

In a range of social learning settings, owing to the information externality, agents conform surprisingly quickly, blocking information aggregation. In other settings (e.g, Smith and Sørensen (2000) and Vives (1997)), the information externality does not lead to complete blockage of learning. But even in such settings, learning is much slower than in the Observable Signals Regime, though not completely blocked (see also Hann-Caruthers, Martynov and Tamuz (2018) and Rosenberg and Vieille (2019)).

Fragility: Fragility is a recurring property in some of the social learning models that we explore in this survey. In the SBM, cascades start at a social belief that is not much more precise than a single private signal. Hence, any comparably informative additional (public or private) signal added to the model could dislodge the cascade—cascades are fragile.[9] In other words, society spontaneously wanders to a position that is highly sensitive to small shocks. Even with more general signal distributions than the SBM, once the social belief is more informative than the most informative signal value, a cascade starts, resulting in fragility.

This contrasts with settings in which social outcomes tend to be insensitive to small shocks (outside particular parameter values that happen to put the system close to a knife edge; Kuran (1989)).

The SBM and the cascades model of BHW provide some surprisingly extreme outcomes: complete blockage of learning along with identical action choices. However, the broader intuition that information externality limits learning suggests that similar but milder possibilities can occur in more general settings: that learning becomes very slow, and that actions become very similar.

Due to its simplicity, the SBM serves as an intuition pump for more complex scenarios. It illustrates which assumptions are crucial and allows us to isolate the effects of changes in assumptions. For example, using a different tie-breaking rule, or allowing for finitely many signal values instead of two qualitatively makes no difference, as shown by BHW. We will likewise see that allowing for endogenous signal acquisition or for endogenous timing of decisions only strengthens the key punchlines. Various other settings also maintain information blockage as in the SBM.

However, there are other modeling variations in which key implications of

[9]We define fragility formally in § 6.

the SBM do not hold. Furthermore, varying assumptions produce interesting new phenomena, such as the endogenous sudden onset of avalanches of activity by many agents triggered by the action of a single agent, eventual learning, or shattering of cascades after their formation. Two insights of the SBM, information externality and herding, are preserved in a very wide spectrum of models, resulting in slow social learning.[10]

These extensions sometimes involve relaxing the SBM's common knowledge assumptions, for example by allowing for privately known preferences, observations and signal precision. In the SBM, agents observe all predecessors and no one else. In more general models of learning in social networks, different agents may observe different subsets of predecessors, and agents may not always know what their predecessors have observed. In models with imperfect rationality, the true economic environment is not common knowledge as agents may systematically misestimate the information sets of predecessors. Exploring different directions for generalization, and isolating the effects of different features of the social learning environment, is a focus of much of the remainder of this survey.

3 Varying the action space

In the SBM, agents chose from one of two actions, and receive a utility that is one if the action matches the state, and zero otherwise. More generally, agents choose an action from some action set, and receive a utility that depends on the action and the state.

The SBM is extreme in having only two actions. At the other extreme, the set of available actions is continuous and, under reasonable assumptions, even a small variation in an agent's belief causes a small shift in action. If so, agents' actions always

[10]An exception is the model of Vives (1997), where there is an information externality but no herding. Since actions are taken in a continuum with full responsiveness to signals, there are no cascades or herding.

depend on their private signals, and therefore reveal their private signals. Hence there is no cascade, and late agents eventually learn the state with near perfection.

Between the binary and continuous extremes, the actions sets that agents choose from are large but finite. For example, firms may choose some technology to adopt, and people choose which sports event to attend, whether to marry one person or another, or how many children to have. Voters must usually choose one from among a set of candidates. Moreover, continuous choice sets are often truncated. For example, the alcohol consumable at dinner is bounded below at zero and above by the limits of human capacity. Furthermore, there is evidence that people sometimes discretize their choice sets, such as restricting their investment orders (price, quantity, or market value) to round numbers (Harris (1991), Kahn, Pennacchi and Sopranzetti (1999), Ikenberry and Weston (2008), Hirshleifer et al. (2019)).

When agents have more available choices, they can potentially attune their actions more finely to their private signals. This can allow later agents to learn more from their predecessors' actions. This in turn suggests that the problem of incorrect cascades may diminish as the action space is enriched. Indeed, one might think that a sufficiently rich action space will make the problems of social learning completely vanish. We will see that this is only sometimes true.

As an example of a simple setting in which agents eventually learn the state, suppose that the two possible states are $\theta = 0$ or 1. Agent I_n chooses an action a_n from the interval $[0, 1]$ and chooses the payoff that minimizes the mean-squared error, i.e., the utility function is $u(\theta, a_n) = -(\theta - a_n)^2$. In this setting, I_n's optimal action is strictly increasing in her beliefs and therefore perfectly reveals her private information. More generally, Ali (2018*b*) defines a notion called *responsiveness*, which loosely speaking means that actions are sensitive even to small differences in beliefs. He shows that responsiveness is a sufficient condition for eventual learning of the true state, for any private signal distribution. Intuitively, if I_n's chosen action is always even slightly responsive to I_n's private signal, subsequent agents can learn I_n's private signal. This results in eventual convergence of beliefs to the true state.

This reasoning may seem to suggest that a continuous action space always eliminates cascades, and results in eventual learning of the state. However, even with a continuous action space, an agent's optimal actions could be unresponsive to the signal. This occurs if there is a particular action that is optimal at more than one belief, in which case it is optimal at an interval of beliefs. For example, this occurs if the action space is $[0, 3/4]$, and the utility is, as above, $u(\theta, a_n) = -(\theta - a_n)^2$. In this case $a_n = 3/4$ is optimal for all beliefs above $3/4$. When the interval at which a single action is optimal is large enough as compared to the informativeness of the signals, no information is revealed about private signals once this region is entered, and a cascade starts. When the set of actions is finite this always happens, and so a cascade starts (and never ends) almost surely (in the spirit of Proposition 1 in BHW). With positive probability, this cascade is incorrect.

Restricting to mean squared error preferences and binary signals, and allowing for multiple states, Lee (1993) obtains a necessary and sufficient condition for fully revealing information cascades. A necessary condition is that the action space not be finite. With a finite action space, choices cannot be responsive; at least one action has the property of being optimal for a range of probability beliefs. This makes actions less informative about agents' private signals. In consequence, incorrect cascades can arise, i.e., there is idiosyncrasy.

Models with a continuum of actions usually have continuous utility functions. An early and prominent exception is the model of Banerjee (1992), in which incorrect cascades occur despite a continuous action space. The unknown state θ is uniformly distributed on $[0, 1]$, and agents choose an action $a \in [0, 1]$. Each agent obtains a payoff of 1 if she chooses action $a = \theta$ and a payoff of zero if $a \neq \theta$. Each agent receives either no signal (and is uninformed) or one signal. An informed agent receives a signal about θ that is either fully revealing or is pure noise (in which case it is uniform on $[0, 1]$). An informed agent does not know which of these two possibilities is the case.

In this model, the payoff from matching the state is substantially greater than from even a slight miss. There is a positive probability that any given predecessor's action a^* matches θ. So when an agent observes that a predecessor chose action a^*, the expected payoff from choosing $a = a^*$ is greater than from choosing $a = a^* + \epsilon$, even when $|\epsilon| > 0$ is arbitrarily small. In consequence, the optimal action is not responsive to small changes in an agent's signal. Thus, early agents may fix upon an incorrect action—an incorrect cascade.

A broader question is whether increasing a finite number of action choices helps prevent incorrect cascades, or at least limits their adverse effects. From a welfare viewpoint, we are now interested not just in whether cascades are incorrect, but in how large the expected mistakes are.

As earlier, suppose that the H state of the SBM corresponds to value $\theta = 1$, and the L state to value $\theta = 0$. As in the SBM, signals are binary, h or ℓ with precision p about θ. However, now suppose further that there are M possible action choices that are evenly spaced between 0 and 1 and always include these two values. (In the SBM, $M = 2$.) So the action set is defined to be $\{0, 1/(M-1), \ldots, (M-2)/(M-1), 1\}$. When $M = 3$, the possible actions are 0, $1/2$ and 1.

Under mean squared error preferences $u(\theta, a_n) = -(\theta - a_n)^2$, each I_n optimally chooses the action a_n that is closest to her inferred probability that the true state is 1. (For example, if the agent's posterior belief about the state is $1/3$, and the choices are $[0, 1/4, 1/2, 3/4, 1]$, the agent chooses $1/4$.) We continue to assume that when an agent is indifferent, she chooses the action that corresponds to her private signal.

In this setting, cascades form with probability 1 (for reasoning similar to the proof of BHW Proposition 1). We measure the social inefficiency as the expected disutility of the late agents, $\lim_{n\to\infty} E_n[(\theta - a_n)^2]$.

The inefficiency is not in general monotonic in the number of actions. In fact, efficiency can decrease when the action set is increased by adding an action without dropping any existing ones. To see this, suppose that the signal strength is $p = 0.7$. When $M = 2$ the set of actions is $\{0, 1\}$ and I_1 follows her signal ($a_1 = 1$ if $s_1 = h$ and is 0 otherwise). Thus, I_2 learns I_1's signal. Adding a third action ($M = 3$) at $1/2$ induces I_1 to choose $1/2$ regardless of signal, as the more extreme actions of 0 or 1 (depending on whether the signal was ℓ or h) would leave her too exposed to the quadratic disutility of action error. Since I_1's action is uninformative, all subsequent agents are in a cascade upon the same choice, $a = 1/2$.[11] More generally, when more actions are introduced, agents may prefer to cascade upon one of the intermediate choices, to the detriment of later agents, rather than using their own signals to take more extreme and informative choices.

Nevertheless, as the number of actions M becomes large the agents cascade upon an action that is likely to be at or near the correct action, and the inefficiency converges to 0. Indeed, even a few choices can significantly improve the sensitivity of action to signal, thereby improving information aggregation (Talley (1999)). For example, when $p = 0.7$, moving from $M = 3$ to $M = 5$ reduces the inefficiency from 0.25 to less than 0.04. On the other hand, when signals are noisy, there can be significant inefficiencies even for surprisingly fine action spaces. For example, if the signal strength $p = 0.51$, an increase from $M = 3$ to the much more refined $M = 49$ does not increase efficiency. Agents continue to cascade immediately on the middle action of $1/2$ (or an action near it), which results in low accuracy of the cascade. Overall, owing to information cascades, even a modest degree of nonresponsiveness of actions to signals (due to the finite number of action choices) can result in substantial social inefficiencies.

In this example, agents have binary signals. However, the same conclusions apply for any signal distribution, as long as signals are weak enough. There is even inefficiency with (weak enough) continuous signals. Intuitively, with sufficiently noisy signals, agents' actions tend to be pushed toward intermediate actions. This leads to

[11] This is not driven by tie breaking and the choice of $1/2$ for the additional action; the same would occur for any number close enough to $1/2$.

frequent cascading at or near the center wherein all agents follow the first agent. What matters for the example is that with a sufficiently noisy signal, the induced belief is never far from $1/2$.

Another property of the analysis that is robust to a change in the signal distribution is that as M becomes larger, information aggregation eventually improves and the disutility of late agents tends to 0. For a given signal distribution, as M becomes large, eventually I_1 is willing to venture outside the intermediate actions, as doing so increases the maximum possible error very little. Information aggregation improves, and inefficiency declines. Similarly, the conclusion that efficiency is nonmonotonic in the number of actions M also applies under more general signal distributions.

4 General Signal Distributions

In the SBM there are only two possible signal values, h or ℓ, but information is often finer-grained. For example, a corporate manager might receive a profitability forecast about a possible investment project that can take one of many possible values.

The natural questions are then: For which signal distributions do information cascades still arise? When do agents eventually herd upon identical actions? What is the effect of the private signal distribution on the probability of settling on the correct action? Understanding these questions requires technical specifics, so this section is more formal than much of this survey.

As in the SBM, we assume a binary state and binary actions, and that I_n takes an action a_n after observing I_1 through I_{n-1}, with the goal of matching the state. We also still assume that private signals s_n are i.i.d. conditional on the state θ, but we now allow a general signal distribution.[12] We still assume that private signals are inconclusive about the state, i.e., the private belief as defined in § 2, $b_n = \mathbb{P}[\theta = H|s_n]$, is in $(0, 1)$. Let the *social belief* be defined as

$$P_n = \mathbb{P}[\theta = H|a_1, \ldots, a_n],$$

the belief held by an outsider who can observe the agents' actions and has no private information. It is convenient to define the *social log-likelihood ratio*

$$R_n = \log \frac{P_n}{1 - P_n},$$

[12]This setting is very flexible. For example, it can capture a situation in which agents have private information about the precision of their own signals. To see this, consider a signal that is comprised of two parts. The first is a symmetric binary signal taking values of either ℓ or h, as in the SBM. The second is a value p that is random and independent of the state, and which gives the precision of the binary signal.

and the *private log-likelihood ratio*[13]

$$r_n = \log \frac{\mathbb{P}[s_n|\theta = H]}{\mathbb{P}[s_n|\theta = L]}.$$

By Bayes' rule, agent I_n's *posterior log-likelihood ratio* for the two states is $R_{n-1} + r_n$:

$$\begin{aligned}
\log \frac{\mathbb{P}[\theta = H|a_1, \dots, a_{n-1}, s_n]}{\mathbb{P}[\theta = L|a_1, \dots, a_{n-1}, s_n]} &= \log \frac{\mathbb{P}[a_1, \dots, a_{n-1}, s_n|\theta = H]}{\mathbb{P}[a_1, \dots, a_{n-1}, s_n|\theta = L]} \cdot \frac{\mathbb{P}[\theta = H]}{\mathbb{P}[\theta = L]} \\
&= \log \frac{\mathbb{P}[a_1, \dots, a_{n-1}|\theta = H]}{\mathbb{P}[a_1, \dots, a_{n-1}|\theta = L]} \cdot \frac{\mathbb{P}[\theta = H]}{\mathbb{P}[\theta = L]} \cdot \frac{\mathbb{P}[s_n|\theta = H]}{\mathbb{P}[s_n|\theta = L]} \\
&= R_{n-1} + r_n.
\end{aligned}$$

Since the utility of an action depends only on the state, the agent's expected utility of an action, given the agent's information, is a monotonic function of the posterior log-likelihood ratio. In consequence, an agent will optimally take action H when this ratio is positive, and action L otherwise.[14]

The social belief P_n (or, equivalently, the public log-likelihood ratio R_n) converges almost surely as the number of agents becomes large, though not necessarily to the truth. This is a general property of any sequence of beliefs of Bayesian agents who collect more information over time. This convergence follows from the Martingale Convergence Theorem (MCT), and the fact that the sequence $P_1, P_2, \dots$ is a *bounded martingale*. More generally, the MCT is a useful tool for analyzing asymptotic outcomes in social learning settings, as it ensures convergence of beliefs under rather mild assumptions.

[13]This definition of the private log-likelihood ratio holds when private signals are discrete. In the case that the conditional private signal distributions have densities, r_n will equal the logarithm of the ratio of the densities.

[14]When $R_{n-1} + r_n = 0$, the agent is indifferent between the two actions. The tie-breaking convention is not important for our conclusions in this section.

4.1 Bounded vs. unbounded private signals

The distribution of private signals, and hence of the private beliefs b_n, is key to understanding whether cascades occur, and whether agents eventually learn the state. If agents may receive arbitrarily accurate signals, then it is intuitive that agents will end up making very good decisions. In contrast, bad outcomes are possible when signals are not arbitrarily accurate, as in the SBM.

The concept of arbitrarily accurate signals is formalized with the notion of unboundedness. We say that private signals are *bounded* if the resulting private belief b_n is not arbitrarily extreme, i.e., there is some $\varepsilon > 0$ such that the belief is supported between ε and $1 - \varepsilon$:

$$b_n > \varepsilon \text{ and } b_n < 1 - \varepsilon \quad \text{almost surely.}$$

We say that a private signal is *unbounded* if, for every $\varepsilon > 0$, the probabilities $\mathbb{P}[b_n < \varepsilon]$ and $\mathbb{P}[b_n > 1 - \varepsilon]$ are both non-zero. With an unbounded signal distribution, signal realization sometimes have arbitrarily high informativeness. Boundedness is a property of the possible beliefs induced by a signal, rather than the range of the signal values per se.

An example of a bounded signal on $[0, 1]$ is one which has density $f_L(s) = 3/2 - s$ in state L and $f_H(s) = 1/2 + s$ in state H. If, instead, the conditional densities in the two states are $f_L(s) = 2 - 2s$ and $f_H(s) = 2s$, then the signal is unbounded. The log-likelihood ratio $\log \frac{f_L(s)}{f_H(s)}$ (and therefore the belief) is bounded in the first case, but can take any value in the real line in the second case.[15]

As shown in Smith and Sørensen (2000), agents eventually learn the state (in a sense to be made precise later) if and only if signals are unbounded. If the set of possible signal values is finite, having unbounded signals is equivalent to having some signal values that reveal the state. We have ruled out signals that reveal the state,

[15] A signal can be neither bounded nor unbounded; for example, it could be the case that the private beliefs can take values arbitrarily close to 1, but are bounded away from 0.

but in fact the possibility of state-revealing signals implies eventual learning of the state, as do unbounded signals more generally. Kartik et al. (2022) point out that when there are three or more states, unbounded signals (appropriately defined for the case of multiple states) are ruled out if the signal distribution has full support and is monotone. The authors obtain sufficient conditions on signals and utility functions that guarantee eventual learning of state.

4.2 Evolution of the Social Belief

We next describe how agents' actions, and thereby the social belief, evolves over time. From the viewpoint of an outsider, the agents' actions form a simple stochastic process with an elegant stationarity feature: An agent's action depends on history only through the social log-likelihood ratio, which captures all the information contained in past actions about the probability of the H state. So following Smith and Sørensen (2000), we denote by

$$\rho(a, R', \theta') = \mathbb{P}[a_n = a | R_{n-1} = R', \theta = \theta'] \tag{1}$$

the probability that agent I_n takes action a when the social log-likelihood ratio is R' and the state is $\theta = \theta'$. The stationarity property is manifested in the fact that the function ρ does not depend on n.

Using ρ, we can define a *deterministic* function ψ that gives the social belief at time n, given the social belief at time $n-1$ and the action at time n, so that $R_n = \psi(R_{n-1}, a_n)$. Since an observer of social information updates her LLR based on the information contained in the latest action, the updated belief satisfies

$$\psi(R_{n-1}, a_n) = R_{n-1} + \log \frac{\rho(a_n, R_{n-1}, H)}{\rho(a_n, R_{n-1}, L)}. \tag{2}$$

This derives from applying Bayes rule to calculate the updated social belief P_n from P_{n-1} and a_n.

4.3 Asymptotic Learning, Cascades, and Limit Cascades

Under cascades, agents sometimes make incorrect decisions even in the long run. To put this in a broader perspective, we now define asymptotic learning—a situation in which the social belief becomes arbitrarily accurate. We then consider the conditions under which asymptotic learning occurs. We also define limit cascades, which are a variant of cascades.

4.3.1 Asymptotic learning

We say that there is *asymptotic learning* if the social belief P_n almost surely tends to 1 with n when the state is H, and to 0 when the state is L. Two other possible definitions for asymptotic learning are the following:

1. The sequence of actions $a_1, a_2, \ldots$ converges almost surely to θ, i.e., all agents from some I_n on take the action that matches the state.

2. The probability that agent I_n takes an action that matches the state tends to one with n.

In this setting both of these are equivalent to asymptotic learning. Intuitively, if beliefs are almost perfectly accurate, so are actions. And for actions to be almost always accurate under the infinite range of possible signal realization sequences, beliefs must also almost always be highly accurate. When asymptotic learning fails, there is idiosyncrasy, as defined informally in § 2.

4.3.2 Information Cascades

Our definition of information cascade in § 2, that the agent takes the same action regardless of her private signal, can be rephrased as follows. An information cascade occurs when, for some agent I_n, the social log-likelihood ratio R_{n-1} is either so high or so low that no private signal can influence the action, i.e., cause I_n's action to depend on I_n's private signal. This occurs if and only if

$$\psi(R_{n-1}, H) = \psi(R_{n-1}, L) = R_{n-1}, \tag{3}$$

i.e., whenever the social belief remains unchanged after observing the action.

4.3.3 Limit Cascades

A *limit cascade* occurs if the agents' actions converge, the limiting action is sometimes incorrect, and each agent chooses either action with positive probability. As with asymptotic learning, the probability that an agent chooses differently from her predecessor decreases quickly enough that, with probability one, from some point on, all agents choose the same action. As with cascades, this action is often incorrect, i.e., the social outcome is idiosyncratic. More formally, in a limit cascade the social log-likelihood ratio R_n tends to a limit that is not $+\infty$ or $-\infty$, but (unlike cascades proper) does not reach that limit in finite time (Smith and Sørensen (2000)); Gale (1996) provides an example of essentially the same phenomenon. Thus, the outside observer's belief converges to an interior point in $[0, 1]$.[16]

If the latest k agents all take the same action, then as k increases, the next agent follows the latest action under a wider range of private signal values. This causes the actions of successive conforming agents to become less and less informative. But since each agent's action always depends on her private signal, actions are never completely uninformative. This contrasts with information cascades, in which actions become completely uninformative. In limit cascades, social learning becomes arbitrarily slow without ever stopping.

Moreover, the informativeness of conforming agents' action drops so quickly that agents do not, in the limit, learn the state. Empirically, the implication is essentially

[16] The limit cascades concept differs from cascades, since a cascade occurs at a point in time, whereas a system is identified as being in a limit cascade only at time infinity.

identical to that of information cascades: society may fix upon an incorrect action forever.

4.3.4 Conditions for the Three Possible Social Learning Outcomes

There are three possibilities for the asymptotic outcome of the process: (i) an information cascade, (ii) a limit cascade, and (iii) neither of the two. These three cases can be equivalently characterized by the limiting behavior of the social log-likelihood ratio R_n, which must converge by the Martingale Convergence Theorem. In case (i) R_n converges in finite time to a finite R_∞. In case (ii) it again converges to a finite R_∞, but not in finite time. And in case (iii) it converges to some $R_\infty \in \{-\infty, \infty\}$. It follows that in this latter case there is asymptotic learning. Smith and Sørensen (2000) describe the relation between signal structures and these three possible outcomes.

In the SBM, the only possible outcome is an information cascade. As shown by BHW, this holds more generally when the set of possible signal values is finite. When signals are bounded but not finite, either cascades or limit cascades can occur. For example, there are always limit cascades when private signals are distributed uniformly on $[0, 1]$ in state L and have density $f(s) = 1/2 + s$ on $[0, 1]$ in state H (Smith and Sørensen (2000)).

When signals are unbounded, there is always asymptotic learning. In other words, almost surely R_n converges to $+\infty$ or $-\infty$ and the social belief converges to either 1 or 0, respectively. Asymptotic learning fails when cascades or limit cascades occur with positive probability. In this case, the social belief converges to some $P_\infty \in (0, 1)$, and the probability that agent I_n chooses correctly tends to $\max\{P_\infty, 1 - P_\infty\} < 1$.[17]

Why do bounded signals block asymptotic learning? Signals are bounded if

[17] For example, when a cascade starts, the social belief reaches P_∞. If $P_\infty > 0.5$, the agent adopts, and this is correct with probability P_∞. Similarly, if $P_\infty < 0.5$, the agent rejects, and this is correct with probability $1 - P_\infty$.

and only if there is some finite M such that the private log-likelihood ratio r_n induced by I_n's private signal is contained in $[-M, M]$. Thus, whenever the social log-likelihood ratio R_{n-1} exceeds M in absolute value, the agent disregards her own signal, since $R_{n-1} + r_n$ must have the same sign as R_{n-1}. It follows that if $|R_{n-1}| > M$, (3) would hold, and a cascade would ensue. Thus it is impossible that $\lim_n R_n = \pm\infty$.

Unbounded signals imply that information cascades and limit cascades are impossible, and asymptotic learning is obtained. To see why, suppose that R_∞, the limit of R_n, is finite, so that a limit cascade or a cascade occurs. Let q_n be the probability that I_n chooses H conditioned only on social information. This probability depends only on R_n, and is strictly between 0 and 1 for any R_n, since signals are unbounded. In the long run since R_n converges to some finite R_∞, q_n will approach some probability $0 < q_\infty < 1$. So an observer who sees the action sequence is essentially receiving an infinite stream of binary signals of approximately constant precision. This observer's beliefs become perfectly accurate, which contradicts the premise that the social log likelihood ratio is bounded.

To sum up, in this setting, there is asymptotic learning if and only if signals are unbounded. When signals are unbounded, no matter how long past agents have been following a mistaken action, there will eventually be an agent with a strong enough signal who will overturn that action. Conversely, when signals are bounded, then the social belief cannot become too accurate, since a cascade would be triggered, blocking further information aggregation.

When there is a cascade, resulting in poor information aggregation, an additional exogenous shock can easily dislodge cascading on the action that was most popular before the shock. So, as in the SBM, outcomes in the setting with general bounded signal distributions can be fragile, a concept developed explicitly in § 6.

A largely unexplored question is the extent to which unbounded signals still succeed in bringing about asymptotic learning when agents are not expected utility maximizers. An exception is Chen (2021), in which agents know the precisions of their own private signals but not of other agents' signals, and are ambiguity averse with respect to this parameter. This can deter agents from breaking a cascade, as agents may fear the worst-case scenario (from the perspective of deviating from the cascade) that predecessors in the cascade are very well-informed. In contrast, when contemplating following one's own signal, an agent is not very fearful of joining the cascade, as the worst that happens is that the agent loses the benefit of the agent's noisy signal. In consequence, even when signals are unbounded, there can be information cascades and no asymptotic learning.

4.3.5 Herds and Speed of Convergence

A *herd* is a realization in which all agents behave alike from some point on. Formally, a herd starts from agent I_n if all later agents take the same action, that is, if $a_m = a_n$ for all $m \geq n$. BHW show that when private signals are finitely supported, herding occurs with probability one. Smith and Sørensen (2000) further show that a herd occurs with probability one more generally for any bounded or unbounded signal structure. It follows that when signals are unbounded, with probability one agents eventually take the correct action. This is somewhat surprising; it happens despite the fact that each agent in the herd has a positive probability of taking either action, i.e., there is herding without cascades. What is even more surprising, as we will see in § 8, is that in other social learning settings there can also be cascades without herding.

So a natural question is: under asymptotic learning, how long does it take for a correct herd to form, and how does the speed of learning depend on the signal distribution?

Asymptotic learning may be much slower than the social optimal rate that would be achieved if private signals were disclosed publicly. In this sense asymptotic learning may be very inefficient. The information externality remains; each agent takes the action that is best for her without regard to the information that her action conveys to others.

Smith and Sørensen ask in particular whether the expected time until a correct herd forms (i.e., until the last time the wrong action is taken) is always infinite. Hann-Caruthers, Martynov and Tamuz (2018) show that this expectation can be either finite or infinite, depending on the tail of the distribution of the private beliefs.

Rosenberg and Vieille (2019) also study how quickly agents converge to the correct action when signals are unbounded. They focus on the expectation of the *first time* that a *correct* action is taken; this is (perhaps not obviously) very closely related to the time at which a correct herd starts. Their very elegant main result is that this is finite conditioned on a given state if and only if $\int \frac{1}{1-F(q)} dq$ is finite, where F is the cumulative distribution function of the private belief in that state. Thus, when private beliefs have very thin tails on the "correct" end—i.e., very low probabilities of extremely informative correct-direction signals—the expected time can be long.

Even with continuous action spaces and unbounded signals, the learning process may still be very slow, as shown by Vives (1993; 1997). In these models, agents only observe a noisy signal about the average actions of predecessors, which relaxes the assumption that the actions of others is common knowledge. Similarly, Chamley (2004*a*), Acemoglu et al. (2011) and Dasaratha and He (2019) provide models with slow convergence. The review of Gale (1996) points out that very slow asymptotic learning, as occurs in several models, may be observationally indistinguishable from complete learning stoppages as in settings with incorrect cascades.

4.3.6 Modeling Considerations

It is largely a matter of convenience whether to model signals as unbounded, so that incorrect cascades never occur, or bounded, so that either cascades or limit cascades occur. There is no way to empirically distinguish a signal distribution that includes values that are extremely rare and highly informative, versus one where such values do not exist at all. So for applications, either modeling approach is equally acceptable. (For a similar perspective, see Gale (1996) and Chamley (2004*b*).)

In many applied contexts, the signal space is in fact finite.[18] In such contexts, we expect to see cascades rather than limit cascades.

Regardless of whether the setting implies cascades or limit cascades, if the setting is modified so that acquiring private signals is even a little costly, then, as discussed in § 7, agents eventually stop acquiring private information. So information aggregation is completely blocked regardless of whether signals are unbounded.

4.4 Heterogeneous Precision and Influencers

Social psychologists report that people imitate the actions of experts. When a sports star uses a particular brand of equipment, it is an expert product endorsement likely to sway observers, as with Roger Federer's endorsement of Wilson tennis rackets and gear. Of course, sports gear often has a non-utilitarian "fashion" element. However, it is plausible that a tennis star knows how to choose an effective racket. Such endorsements may be less compelling for products that are unrelated to the endorser's primary domain of expertise. However, even outside this domain, observers may still view the decisions of an exceptionally successful individual as reflecting knowledge about what choices are most effective.

As the tennis example illustrates, some agents may have systematically more accurate signals than others. This raises the questions of whether small differences in accuracy can make big differences for social outcomes, whether increasing the accuracy of some agents necessarily improves social outcomes, and whether agents with higher

[18]For example, people often obtain information signals through casual conversation with limited nuance. ("Is that movie worth seeing?" "Yeah.") There is evidence from psychology that there are minimum distinguishable gradations in sensory cues. People often obtain information from experiments with a small number of possible outcomes, such as learning whether a job applicant is or is not a high school graduate. Course grades are discrete (such as A through F), as is learning whether a new acquaintance is married or unmarried, how many children the individual has, or who the individual voted for in the last election.

accuracy should be placed earlier or later in the decision queue (Ottaviani and Sørensen (2001) study the latter question in a reputational learning setting).

Consider a variation of the SBM in which some agents, whom we call *influencers*, have more precise private signals than other agents. Depending on their location in the decision queue, influencers can trigger immediate cascades. So even a small advantage in signal precision can have a large effect. For example, if we alter the SBM so that Ann has a signal precision $p' > p$, then even if p' is close to p, Bob will defer to her, as will all later agents. The resulting cascade has precision p'. In contrast, if $p' = p$, a cascade occurs only after two identical actions, which makes use of at least two signals, and therefore is more accurate.

The outcome depends crucially on the order of moves. If the high-precision agent, Ann, were second instead of first, there would be no immediate cascade. Placing an influencer later allows the actions of early agents to remain informative, which improves the accuracy of the cascade that ultimately ensues (see also § 12.1). These effects of influencers apply also in settings with with more general signal distributions, even though an influencer may not immediately trigger a cascade.

The drawback of leading off with the better-informed has not been lost on designers of judicial institutions. According to the Talmud, judges in the Sanhedrin (the ancient Hebrew high court) voted on cases in inverse order of seniority. Similar voting orders continue in some of today's courts (e.g., those in the U.S. Navy). Such strategic ordering can reduce the undue influence of older (and presumably wiser) judges.

4.5 Partial Cascades

The logic of information cascades is that with a coarse action space, an agent may choose an action independently of her private signal. In consequence, her action does not add to the pool of social information, and asymptotic learning fails. Furthermore, even in settings where there is asymptotic learning, the improvement of social information tends to be delayed.

This logic is a special case of a more general point: that a coarse action space reduces the informativeness of an agent's action. This effect is in part mechanical. For example, if there were just one possible action, taking that action would convey no information. However, information loss is exacerbated by information externality; self-interested agents have no incentive to convey information to later agents.[19]

To capture this more general point, Lee (1998) defines a weaker notion of information cascade. A *partial cascade* is a situation where an agent takes the same action for multiple signal values. (This terminology is due to Brunnermeier (2001); Lee uses the term "cascades" for this concept.) An information cascade proper is the special case in which an agent takes one action for all the possible signal values. Partial cascades occur trivially when there are more possible signal values than actions. But even when the action space is rich, partial cascades (and cascades proper) can occur.

Lee applies this notion to a model of information blockage and stock market crashes (as discussed in § 12). In the model of Hirshleifer and Welch (2002), partial cascades result in either excessive maintenance of early actions ('inertia") or excessive action shifts ("impulsiveness"). Overall, there is a wide array of settings in which partial cascades hinder information aggregation.

5 Endogenous timing of actions

When faced with an irreversible decision, people often have an option to either act immediately, or to defer their decision to a time in which more information will be at their disposal. Examples include purchasing versus deferring a product, or undertaking versus deferring an investment project.

[19] We mainly focus on coarseness in conjunction with information externalities. Coarseness is much less of a problem with altruistic agents, as they could choose their actions strategically to convey their private information to others (see Cheng, Hann-Caruthers and Tamuz (2018)).

Consider a modification of the SBM in which, at any given date, any agent is free to choose among three options: adopt, reject, or delay. Adopting or rejecting is irrevocable. In this setting, there is no exogenous sequencing in the order of moves.

The benefit to delay is that an agent can glean information by observing the actions of others. Thus, delay generates option value. The cost of delay could take the form of deferral of project net benefits, or of ongoing expenditures needed to maintain the option to adopt. Since acting early confers a positive information externality upon other agents, in equilibrium there can be excessive delay.

This equilibrium outcome when agents have an incentive to delay is seen mostly simply when agents have heterogeneous signal precision. We consider this in the next subsection. This setting provides insight into a wider set of models considered in the remainder of this section.

5.1 Delay with Heterogeneous Signal Precision

Suppose that agents differ in the precisions of their binary private signals (extending the SBM to agent-specific values of p that are either commonly or privately known). At a given point in time, an agent can *act* by choosing project H or project L; or alternatively, can delay. As in the SBM, project H is optimal in state H, and L is optimal in state L. As discussed in BHW (p. 1002), high-precision agents have less to gain from waiting to see the actions of others—in the extreme, a perfectly-informed I_i ($p_i = 1$) has nothing to gain from waiting. So we focus on equilibria in which, among agents with h signals, those with higher precision adopt earlier than those with lower precision.[20]

[20]Our focus on equilibria in which agents choose different timing implicitly requires that time periods be short relative to differences in possible precision. This ensures that an agent with lower precision would prefer to wait one period and learn from a higher-precision agent rather than act simultaneously. To avoid technicalities, we suppose that agents have precisions drawn without replacement from a discrete distribution, such that no two agents have the same precision. This ensures that, if time periods are short enough, different agents act at different time periods.

Suppose first that each agent's signal accuracy is known to all. The binary signal of the agent with the highest precision dominates the signals of all other agents, as in the influencers example discussed in § 4.4. Once the agent with highest precision has acted (H or L), all remaining agents are in a cascade on the selected action. Intuitively, the agent with second-highest precision (the one with the next least gain from delay) acts immediately rather than waiting to learn from others. Since this action is uninformative, for similar reasons, so do all others.

Now suppose instead that precisions are only privately known. In equilibrium, agents can infer the signal accuracy of other agents from time elapsed without action. In the continuous-time model of Zhang (1997), this results in an equilibrium in which delay fully reveals precision. Each agent has a critical maximum delay period, after which, if there are no actions by others, she becomes the first to act. The higher the precision, the shorter the critical interval. Again, all agents wait until the highest-precision agent acts. At that point, all other agents immediately act in an investment cascade.[21]

In this model, cascades are *explosive* in the sense that there is an initial time period during which all agents delay, and then, once the highest-precision agent acts, others immediately follow. Furthermore, since the cascade is based solely on one agent's signal, actions are also highly idiosyncratic.

In the real options model of Grenadier (1999), the underlying asset value evolves as a geometric Brownian motion with an unknown parameter. Investors differ in the accuracy of their private signals about the value of the unknown parameter. If two high-accuracy investors with positive signals invest, all other agents are in an information cascade, and also invest. Thus the broad insight from Zhang (1997) holds in Grenadier's setting as well—that agents with more accurate signals have a stronger incentive to act first, and that once this occurs, low precision agents have an incentive

[21] The insight that the incentive to free ride on the information of others by delaying also holds in a setting with continuous signals and actions, and where information acquisition may be endogenous Aghamolla and Hashimoto (2020).

to mimic, triggering a flurry of activity.

The fact that less accurate agents have a greater benefit to delay suggests that when decisions are observable to others, there can be a strategic advantage to acquiring less information. Consistent with this, even if information acquisition is costless, agents may choose to remain imperfectly informed, as this encourages other agents to act earlier Aghamolla (2018). Moreover, even in nonstrategic settings, agents may acquire too little information, since there are externalities in information acquisition (§ 7).

5.2 Adopt versus Delay as an Indicator of Degree of Optimism

In the setting we have discussed, delay is informative about agents' precisions, but not about whether their signals favor project H or L. However, suppose now that there is only a single project, and that the decision each period is whether to adopt it or to delay. As before, adopting the project is irrevocable. Then delay can be an indicator that the project is unattractive.

The consequences of this are seen most simply in a two-period model in which agents have identical precisions. Typically, the incentive to delay in such a setting rules out symmetric pure strategy equilibria. To see why, suppose that there were such an equilibrium in which all agents with h signals adopted immediately. Then their actions would accurately reveal the state. The value of defecting by waiting one period for more social information would be very high, breaking the equilibrium. Consider instead a proposed equilibrium in which all agents with h signals delay one period. Since delay is costly, it would pay to defect from the equilibrium by acting immediately rather than acting one period later, since no information is obtained by waiting.

Wang (2017) focuses on asymmetric pure strategy equilibria in which agents with identical precisions delay different amounts of time. When agents are patient enough, the initial prescription of the equilibrium strategies is that in each period a designated agent invests if her signal is high and delays if her signal is low, and all other agents delay. This prescription for the active agents endogenously generates sequential choice that is similar to having an exogenous order of moves as in the SBM. Once sufficient social information accumulates, all agents who have not yet invested cascade, meaning they either invest forever or delay forever. So even when timing is endogenous, in equilibrium action choices are very similar to those in the SBM.

Turning to mixed strategy equilibria, Chamley and Gale (1994) derive randomization in delay in a setting in which a subset of agents randomly receives the option to invest—an option that can be exercised at any time. In practice firms sometimes do not immediately have the resources or technology available to undertake an investment project, so that delay results from capacity rather than choice. In the model, more agents receive the option to invest in better states, so receiving the option is a favorable indicator about state. This is the only signal that agents receive. If many agents receive positive signals, their decisions to invest can suddenly encourage others to follow. So there can be sudden investment booms. However, by the same token, delay by many agents conveys adverse information to others, resulting in an investment bust. In the unique symmetric equilibrium, agents with favorable signals (implicit in receiving the option) delay with positive probability in order to gain information by seeing how many others invest.[22] The information implicit in prolonged delay by many agents can cause additional investment to cease. Also, owing to the externality that agents benefit from waiting to observe what other agents do, there tends to be excessive deferral of investment.

In contrast, if all agents have the option to invest, and each agent receives a direct private signal about the state, there is no bias towards underinvestment relative to full information aggregation, as shown by Chamley (2004*a*). At any date, the number of agents who have already invested is a positive indicator about the state. There are

[22] The idea that information externalities result in stochastic delay is in the spirit of Hendricks and Kovenock (1989), who examine an experimentation setting in which two firms with private information decide how soon to drill for oil, where drilling causes the arrival of public information about the payoff outcome.

multiple equilibria with different thresholds for the belief about the state above which agents invest. In an equilibrium with a higher threshold, the information conveyed by the decision to invest is stronger. Thus, when the threshold signal value is higher, the marginal agent has a higher informational benefit of waiting, which compensates for the higher cost of delaying longer. This supports the equilibrium.

Overall, these models of timing decision reveal that there can be either inefficient delay, or a rush to invest even in unprofitable projects. This suggests that social learning may generate shifts in investment activity that are reminiscent of observed industry-wide booms and busts, or macroeconomic fluctuations.[23] Such shifts occur within equilibria in the Wang (2017) and Chamley and Gale (1994) models, and might be viewed as occurring occasionally across multiple equilibria in Chamley (2004*a*).

In practice, sometimes firms can repeatedly adopt and terminate projects. Caplin and Leahy (1994) analyze information cascades in project decisions when firms can receive multiple private signals over time and can see the actions previously taken by other firms. After an uneventful period of delay, there can be sudden crashes in which many firms terminate their projects at about the same time. Towards the end, with enough signals, firms essentially know the value, and take the correct action. Thus, incorrect cascades do not occur.

As we have discussed, two interesting features that can arise in settings with endogenous timing (Chamley and Gale (1994) and Zhang (1997)) are that (i) agents take similar actions, as in models of information cascades, and (ii) actions tend to be clustered in time. Gul and Lundholm (1995) provide a setting in which social learning and the option to delay can result in time clustering. There are continuous actions and signals, making actions responsive to signals, so that cascades on action choice do not occur. In their setting there is time clustering of the last two agents in a sequence, because once one of them acts, the other does not gain any further information from

[23]Models of delay in which agents can observe the payoffs of predecessors (Caplin and Leahy (1993), Wagner (2018), Aghamolla and Guttman (2021)) have also been applied to such phenomena.

delay. However, their setting does not explain time clustering by large numbers of agents.

6 Observability Assumptions

What an agent learns from others, and whether the agent subsequently falls into a cascade, depends on what the agent can observe about past history. And what can be observed about history is, in practice, highly context-dependent. Greater observability can take the form of observation of payoffs, or even private signals, not just actions. Observations can also be noisy or limited, as with observation of only a specified subset of agents, a random sample, or a count of adoptions or other aggregate statistic. Observability can be asymmetric (as with greater observation of adoptions or of higher payoffs). Furthermore, there can be meta-uncertainty, wherein agents are not certain whom their predecessors have observed.

Some key questions, under alternative observability regimes, are whether incorrect cascades must eventually be dislodged, whether there is asymptotic learning, and whether social outcomes are fragile with respect to the arrival of new public information. Another key question is whether greater observability increases welfare. We first discuss these issues in the context of rational settings, and then turn to imperfect rationality.

6.1 Rational Models

There are many different combinations of observability assumptions in the literature. Rather than systematically discussing the proliferation of different possible assumptions, we focus on some general themes. We organize the discussion in terms of these themes rather than by model assumption.

Recall that in § 4 we defined a herd as starting at agent I_i when all later agents do the same thing as agent I_i.

Theme 1. When there is sufficient observation of past actions (and possibly additional social information, such as past payoffs), the probability that a herd eventually starts is 1 in many models.

As a benchmark, if there is no social observation, agents act based on their own private signals, and there is no herd. However, with enough social observation of actions and perhaps payoffs, more and more information is revealed, at least until agents eventually tilt toward one action. This intuition requires reasonably good observation of past actions. For example, insufficient social observation precludes herding in Celen and Kariv (2004*a*) (which we discuss again in § 11), where each agent observes the action of the immediate predecessor only. There are many models with sufficient social observation, including those of Banerjee (1992), BHW, Smith and Sørensen (2000), Cao, Han and Hirshleifer (2011), Hirshleifer and Plotkin (2021), and models with costly information acquisition discussed in § 7. In some of these models the probability that a herd eventually forms is less than one.

We discuss the next two themes together.

Theme 2. When there is sufficient observation of past actions or payoffs (and possibly other social information), herds can be incorrect. Specifically, incorrect information cascades can occur, and can last forever with strictly positive probability. So in general, asymptotic learning may not occur.

Theme 3. If there are private signals, and the payoffs to predecessors' actions are observable, information cascades can cause insufficient *exploration* as well as poor information aggregation.

A fundamental trade-off in an individual decision-making setting is between taking actions that generate new information that is relevant for an agent's future actions, versus taking the myopically best action based on current beliefs. This is known as the *exploration/exploitation trade-off*. In a social learning setting, even if each agent takes only one action, from a social welfare perspective there is still a tension between the individual benefit of exploiting existing information versus exploring in the sense of generating new information that is helpful for later agents.

Consider now a deviation from the SBM in which previous payoffs are observable but are stochastic given the state. In such a setting, actions may differ in the usefulness of the payoff information that they generate. In consequence, there are two types of information externalities. The first, just as in the SBM, is that in choosing an action, an agent does not take into account that her choice of action conveys information about her private signal to later agents. This externality affects the *aggregation* of private signals.

The second type of information externality is in the *generation* of new information. Agents do not take into account the benefit to later agents of observing the payoffs derived from the chosen action. This is an externality in exploration (see, e.g., Rob (1991)). Owing to this externality, in a social multi-arm bandit settings with no private information, asymptotic learning fails (Bolton and Harris (1999)). An analysis with quasi-Bayesian agents is provided by Bala and Goyal (1998).

In principle, with many agents, there are enough private signals for information aggregation alone to induce asymptotic learning. Similarly, exploration alone could generate enough payoff information to pin down the realized state perfectly. These facts raise the hope that when agents can socially acquire both types of information there would be asymptotic learning. After all, observation of predecessors' payoffs can sometimes dislodge an incorrect cascade, thereby potentially resulting in new trials and payoff observations of both choice options.

Nevertheless, it turns out that even when payoffs are observed, there can be a strictly positive probability that an incorrect cascade forms and lasts forever, i.e., outcomes are idiosyncratic. As in the SBM, once a sufficient predominance of evidence from past actions and payoffs favors one action, agents start to take that action even when their own private signals oppose it. The basic logic of cascades still applies; additional private information is no longer impounded in actions. So this cascade may be incorrect. In contrast with the SBM, such a cascade may be broken by the arrival of payoff information. But asymptotic learning is not assured.

To see this, consider a setting in which all past payoffs as well as actions are observable (Cao, Han and Hirshleifer (2011)). Suppose that the payoffs to action a are either 1 or -1, and to action b are either 2 or -2. There are four equally likely states: uu, ud, du and dd, where the first entry indicates a high (u) or low (d) payoff to action a and the second entry indicates a high or low payoff to action b.

Once an action is taken, its payoff is known to all, an assumption that is highly favorable to effective social learning. Nevertheless, a problem of inadequate experimentation remains. If private signals and payoff information about action a are initially favorable, whereas the prior beliefs about b are not very favorable, society can lock into a, for an expected payoff close to 1 under the belief that ud is likely, without ever trying b, whose payoff in state uu of 2 is even higher. So the ability to socially acquire both types of information (about private signals and about payoffs) does not solve the information externality problem.

More generally, if payoffs are stochastic even conditional upon the state (or observation of payoffs is noisy), there is still a strictly positively probability that a given agent will cascade upon an incorrect action. Furthermore, agents sometimes lock into an incorrect action forever even after having tried both alternatives any finite number of times.

A possible interpretation of the payoff signal is that it is an online review posted by an agent who has adopted. In this application, Le, Subramanian and Berry (2016) show that the probability of an incorrect cascade can increase with the precision of the review. A more accurate review that is favorable to one option can prematurely deter useful exploration of the other option. Likewise, in a model of social learning from online reviews (in which agents have no private information about the state), Acemoglu et al. (2019) find that providing more information does not always lead to better outcomes. When customer heterogeneity is sufficiently high, Ifrach et al. (2019) find that there is asymptotic learning, as there is always a chance of purchase strictly between zero and one, regardless of the social belief.

Alternatively, the additional social information that agents obtain may be about the private signals of predecessors through conversation. If *all* past private signals were observed, then trivially agents would converge to the correct action. In reality people do sometimes discuss reasons for their actions, but they often do not pass on the full set of reasons that they have acquired from others. It is not hard to provide an example with limited communication of private signals in which incorrect cascades form and, with positive probability, last forever. In the Online Appendix, § A.1.3, we provide an extension of the SBM where for all $n > 1$, I_n observes the private signal of I_{n-1}, and incorrect cascades still occur and last forever (see Cao and Hirshleifer (1997)).[24]

Intuitively, seeing the predecessor's private signal is much like directly observing an extra private signal oneself, which in turn is much like seeing a more precise signal. In the SBM, a more precise private signal tends to make decisions more accurate, but the probability that the long-run cascade is incorrect is still strictly positive. So there is reason to expect the same when the predecessor's signal is observed.

In particular, the logic of information cascades applies; a point is reached when the information implicit in past actions overwhelms the bundle of the agent's own signal and the predecessor's signal. Such a preponderance of evidence can still be far from conclusive.

Theme 4. Social outcomes are often fragile with respect to the arrival of modest

[24] An agent might communicate a sufficient statistic (such as the agent's belief) instead of one or a few private signals. In practice, this sometimes occurs, but people also sometimes seem to convey one or two specific reasons rather than conveying an overall degree of belief.

new public information.

The occurrence of small shocks, such as the arrival of even modest public information, can easily dislodge a cascade. Each agent knows that any cascade is based upon information that is only slightly more accurate than the agent's own private signal. Thus, as emphasized by BHW, a key prediction is that even long-standing cascades are fragile with respect to small shocks.

Definition 2. *A cascade is* fragile *if a hypothetical one-time public disclosure of a signal with a distribution that is identical to that of the private signal possessed by a single agent would, with positive probability, break the cascade, i.e., causes the next agent's action to depend on that agent's signal.*

For example, in the SBM, once a cascade starts, it remains fragile for all remaining agents.

Fragility is a general concept that could be defined in terms of different kinds of shocks to the system. For example, instead of the arrival of a public signal, the shock could be the arrival of a better-informed agent, or the arrival of an agent whose preferences are known to differ from predecessors. In each of these cases, a long-standing cascade can easily be dislodged.

Owing to information cascades, there is a systematic, spontaneous tendency for the system to move to a position of precarious stability. This is much like the way the hero's car in an action movie chase scene always ends up teetering at the edge of a cliff. In contrast, equilibrium is much more stable in models in which there are sanctions upon deviants or disutility from nonconformity (Kuran (1987)).

For the next observation, consider a setting with a general private signal distribution, two states and two actions, and simple sequential observation wherein each agent observes all predecessors. We allow here for a wide range of other model features, such as costly acquisition of private signals and endogenous order of moves.

Theme 5. The Two-Signal Principle: The contribution of private signals to the social belief has information content of at most two reinforcing occurrences of the most informative private signal value.

Recall from § 2.1 that when the state is binary, the precision of a possible belief r about the state θ is defined as $|r - 0.5|$.

To understand the two-signal principle, consider first a setting in which the prior belief is symmetric, and in which there is no arrival of public information signals. With no public signals, the conclusion is simpler: the phrase "contribution of private signals to the social belief" in Theme 5 can be replaced with "the social belief."

To see why, observe that there is no way for an agent I_n's action to reflect more than two maximally informative private signal realizations unless at least one earlier agent's action reflects more than one such realization. Consider any point in the sequence in which more than one maximally informative private signal realizations is incorporated into the social belief as observed by I_n. Then I_n will follow the action implied by the social belief regardless of I_n's private signal, i.e., I_n is in a cascade. It follows that I_n's action is not informative, which means that it does not increase the precision of the social belief. So I_{n+1} and all later agents are also in a cascade, and also do not increase the precision of the social belief. In other words, the zone of social beliefs that incorporate between one and two private signals is impassable.

The premise of this argument excludes a common feature of many social learning models: observation of payoff outcomes. Nor does it allow for the arrival of other kinds of public signals. However, with only slight modification, the reasoning above can address the possibility of asymmetric priors and/or the arrival of additional public information. Let us now allow for these.

Just as in the reasoning above, there could be a social belief faced by I_n being $q \neq 1/2$. If this social belief does not start a cascade, then as reasoned above we can get at most only two maximally informative signals in the same direction before starting a cascade. The only other possibility is that public information disclosures have generated a social belief that has high enough precision to start a cascade. If so, we are trivially in the impassable zone, and no private information ever again contributes to the social belief.

The social belief could potentially be very precise. But the above reasoning makes clear that such a possibility would derive directly from publicly arriving information, not from aggregation of private information. It is still the case that less than two private signals are aggregated into the social belief.

Theme 6. The Principle of Countervailing Adjustment: Seemingly favorable shifts in information availability do not necessarily improve average decisions or welfare.

The direct positive effect of more information tends to be opposed by the tendency of agents to disregard their private signals sooner, to the detriment of later agents. We call this the *principle of countervailing adjustment*. An example is the presence of an agent with slightly better private information (an "influencer"). As discussed in § 4.4, increasing the precision of one agent can reduce average welfare by causing later agents to fall more readily into a incorrect cascade. Wu (2021) applies the principle to the formation of cascades in two decision queues.

One consequence of the principle of countervailing adjustment is that disclosure can reduce average welfare. This contrasts with settings with no social interaction, wherein an extra signal always makes an agent weakly better off.

In the influencer model discussed in § 4.4, I_1 has a slightly more accurate signal than later agents, leading to reduced average welfare. Suppose instead that all private signals have identical precisions, and that the slightly more accurate signal is a *public* disclosure made at date 0. Now agent I_1 is the first in the cascade, and all agents have lower expected utility than in the basic setting, as all now effectively act based upon just a single signal (the public disclosure). Of course, a sufficiently accurate early signal or public disclosure can improve the social outcome. For example, if I_1 has perfect information, the cascade is always correct.

More generally, a shift in information regime that might seem to make agents better informed (such as higher signal precision of early agents, greater observability of others, or greater publicly available information) can reduce the average decision accuracy in the long run and can reduce average welfare. An example of this is the possible deleterious effect of increasing the precision of publicly posted information discussed by Le, Subramanian and Berry (2016) above. Intuitively, some variation in the setting directly that makes early agents better informed promotes cascading by observers, to the possible detriment of even later agents.[25]

Theme 7. When each agent observes only a random sample of past actions, incorrect information cascades can occur, and may last forever. So asymptotic learning does not necessarily occur.

The conditions for this theme are potentially compatible with those of Themes 1 and 2 on herding, so there are settings where the conclusions of both hold, i.e., there can be incorrect cascades that last forever; which is compatible with, and indeed implies, herding. To understand this theme, consider a sequential setting with random sampling of past actions and with no information about the order of past actions (see Smith and Sørensen (2020)). With bounded signals, information aggregation tends to be self-limiting, because more informative actions tend to encourage cascading upon the preponderance of actions in the agent's observation sample. Whenever such cascading occurs, the agent's signal is not incorporated into the action history.

Suppose, for example, that past agents' actions were to become so accurate that even a single sampled H action would be sufficient to overwhelm the most extreme possible opposing private signal value. Then an observer will sometimes be in a cascade

[25]However, in general a shift in model structure can have different types of effects on the quality of information aggregation. For example, a shift in model structure in some cases affects the critical value for a cascade to occur, which can either increase or decrease the ultimate amount of information impounded in the action history. Furthermore, even within the SBM, if each agent's private signal became more precise, cascades will tend to be more informative.

upon the predominant action in the agent's sample (e.g., in a sample of size k in which all observations in the sample are of the same action).

This reasoning suggests that owing to the possibility of information cascades, learning with random sampling may be quite slow. Indeed, a stronger claim is true: so long as all agents observe a sample size of at least 1, asymptotic learning fails (Smith and Sørensen (2020)).

To see why, consider the case of a sample size of $N = 1$, where private signals are symmetric and binary. The departure from the SBM is that each agent observes the action of one randomly selected predecessor instead of all predecessors. Suppose that a point is reached where for some agent I_n this random observation is more informative than a single private signal. (If this never happens, of course asymptotic learning fails.) Then agent I_n would be in an information cascade, so I_n's action would be exactly as informative as a sample of one action from among I_n's predecessors. In consequence, the sample observed by I_{n+1} is also more informative than I_{n+1}'s signal, so I_{n+1} would also be in a cascade. A similar argument holds for all later agents, so information stops accumulating. Consequently, asymptotic learning does not occur. This failure is similar to the fashion leader version of the SBM in Subsection 4.4, in which information stops accumulating once an action is observed with precision greater than an agent's private signal. A similar intuition also applies to the random sampling model of Banerjee and Fudenberg (2004), in which it is possible that past payoffs as well as actions are observed.[26]

A further interesting implication of the sequential sampling setting of (Smith and Sørensen (2020)) is that assuming each agent does not observe too many predecessors, once different agents take different actions, agents never herd. This is because there is always a chance that an agent observes a set of predecessors who did not

[26] Monzón and Rapp (2014) consider a sampling setting in which agents also do not know their own positions in the decision queue. In this setting, under a stationarity assumption on sampling rules, incorrect cascades can last forever.

follow the currently-predominant action (Smith and Sørensen (2020)). The chance of observing such a deviant action does not decline rapidly enough with n to bring about herding. Similarly, if payoff information is also observed and if an early agent adopted a popular action and experienced low payoffs from doing so, there is always a chance that this agent is later observed, causing a later agent to deviate from the popular action.

If agents observe samples of payoff outcomes but not the past actions that led to these outcomes, it is again possible that agents do not converge to the same action. The need to simultaneously draw inferences about what actions predecessors have taken and the performance of those actions can confound inferences. Wolitzky (2018) considers a setting with two actions: action R(isky) has a probability of success that depends on state, and action S(afe) has a fixed probability of success that is state independent. If action R potentially generates a higher state-contingent probability of success than action S, then outcomes become close to efficient if the size of samples becomes arbitrarily large. However, when action R always has lower probability of success (but has lower cost) than action S, then even for very large samples, there is not asymptotic learning.

Theme 8. Reject cascades can occur even when agents observe the aggregate number of adopts, but do not observe rejects.

When only aggregate adoption counts are observed, sequencing information is lost. In general this induces loss of two types of information. First, an agent does not know *the order* in which past actions were taken. Second, an agent does not know *how many* predecessors have acted—i.e., agents do not know their own positions in the queue. This occurs when an agent does not observe all past actions, one example being when an agent observes only adopts, not rejects. For example, it is not hard to obtain information about how many Teslas have been sold, but we do not observe how many people considered Teslas but opted not to buy.

In the setting of Guarino, Harmgart and Huck (2011), observation is asymmetric (we will refer to this as only observing past adopts, not rejects), there is a finite number of agents, agents cannot see the order of predecessors' actions, and they have no direct information about their own positions in the queue (though they can draw inferences about this from their observations of past adopts). Since all that an agent observes is how many predecessors adopted, there is no way for a cascade on reject to get started. (If it could, then even I_1 would reject, since I_1 does not see any past adopts nor does I_1 know that no agent preceded her. In consequence, all agents would always reject, which is not consistent with equilibrium.) So the possibility of cascading is limited to just one action, and indeed, with a large (finite) population, such a cascade occurs with a probability that approaches one regardless of state.

In sharp contrast, when agents do have some idea about their own positions in the queue (based, for example, on observation of own-arrival-time), cascades on either action can occur. This is because an agent who is probably late in the queue and who observes few adopts infers that others probably arrived earlier and chose to reject (Herrera and Horner (2013)).

Theme 9. Contrarian actions can reveal that an agent has high precision.

Consider a setting like the influencer model of § 4.4 except that agents have private information about the precisions of their signals. Then the decision of an agent to deviate from a cascade indicates that the agent has high precision. This can potentially cause subsequent agents to follow contrarians.

In this setting an observer knows that the minority choice was made in opposition to predecessors, which is indicative of strong private information. What is more surprising, as shown by Callander and Hörner (2009), is that agents who only observe a count of past adopts and rejects sometimes act in opposition to the majority of the actions they observe.

To see this, consider the SBM, except that I_3 has a conclusive signal, whereas I_1, and I_2 have very noisy signals. We can think of I_3 as a "local" who knows whether a restaurant is good. Even if I_1 and I_2 (low precision "tourists") adopt, if I_3 rejects, and

I_4 understands this structure, clearly I_4 will imitate I_3.

Furthermore, even if I_4 is uncertain about whether I_3 is a local, I_4 can infer this from the fact that I_3 rejected after two adopts. The very fact that I_3 came late and was in the minority (is a "contrarian") is an indicator of his high precision.

What if I_4 does not know the order of moves, only that one predecessor was in the minority (two adopts and one reject) and that (for simplicity) there was exactly one local? If I_1 were local, then the two tourists, I_2 and I_3, would have imitated I_1 in the hope that she is local, which would have generated a unanimous choice. So this possibility is ruled out. If I_2 is the local, there is a 50% chance that I_3, a tourist who observed one adopt and one reject, chose an action that matched that of I_1, so that the local is in a minority. If, instead, I_3 is local, then I_2, being a tourist, would have imitated I_1; thus, the two adopts are by I_1 and I_2 and the reject by I_3, the local. So overall, the preponderance of evidence that two adopts and one reject conveys to I_4 is that the local rejected. Being a contrarian can be an indicator of being well-informed, so if I_4 is a tourist, I_4 optimally follows the minority.

6.2 Models with Imperfect Rationality

In models of individual decision making, irrationality makes an agent worse off. In contrast, in a social learning setting, irrationality can make most agents better off, because mistaken actions are sometimes more informative than correct ones to later observers. So psychological bias can help remedy information externalities, resulting in more accurate social beliefs. This is an example of the general phenomenon that irrationality can make interacting agents better off (Kreps et al. (1982)).

Within a rational setting, BHW and Banerjee (1992) point out that there is a benefit to quarantining early agents so that some make decisions without observing others. We call such agents *sacrificial lambs*. The advantage of having nonsocial agents is that their actions depend on their own signals. This makes their actions more informative to later observers. Similarly, psychological biases can cause agents to use their own signals instead of imitating others, thereby improving learning and welfare.

6.2.1 Overconfidence

Even if an agent is not in an information quarantine, the agent may follow her own private information because she is overconfident about its precision. Such overconfident agents can break incorrect cascades, improving long-run learning. In Bernardo and Welch (2001), occasional overconfident "entrepreneurs" overestimate the precisions of their own signals. This can cause them to make greater use of their own signals instead of following the actions of predecessors in an information cascade. So overconfidence can improve learning and outcomes for later agents.

Suppose, for example, that everyone's private signal has the same precision, known to all, except that I_1 to I_{10} each substantially overestimates their own precision. Then each acts based solely upon the agent's own signal, so the first 10 signals are revealed through their actions. In consequence the expected welfare of all later agents is improved.

A possible direction for future research is understanding the conditions under which, in contrast with Bernardo and Welch (2001), excessive overconfidence harms social learning instead of helping. For example, in an extreme case, if overconfidence were growing rapidly with later agents, all agents would act based only on their private signals, and there would be no information aggregation.

6.2.2 Neglect of Social Observation by Predecessors

Various other psychological biases can also influence social learning. An important one is that agents may neglect the fact that others are making social observations. This induces *correlation neglect* (also known as *persuasion bias*), the phenomenon that people sometimes treat information they derive from others as independent even if there is commonality in the sources of this information—a type of double-counting (Enke and Zimmermann (2019)). Such neglect is a natural consequence of limited attention and cognitive processing power.

Since inferences about observation of others can require extensive computation, it is plausible that agents update beliefs heuristically. To update rationally, an agent needs to adjust for the fact that the information in an observed action depends on whom the actor in turn is able to observe (see, e.g., Acemoglu et al. (2011)). In reality people typically do not adjust appropriately.

As with overconfidence, correlation neglect can cause agents to make greater use of their own private signals instead of cascading, thereby improving welfare. To see why, consider an agent who observes just one predecessor, and who mistakenly believe that this predecessor is not observing others. Owing to this mistake, the agent underestimates the informativeness of the predecessor's action. This makes the agent more inclined to rely on her own signal, which makes her own action more informative to later agents.

However, correlation neglect can also make agents more prone to cascading. To see why, consider an agent who observes multiple predecessors who happen to choose the same action, and mistakenly believes that all these predecessors have acted independently. Some of these predecessors may be in a cascade, making their actions uninformative. The mistaken belief that these actions are all informative can cause the agent to imitate predecessors and join a cascade instead of following the agent's own signal. This can happen even if the agent has higher private signal precision than do predecessors, so that if the agent were rational the agent would not be in a cascade.

In a social network, heavily connected agents provide correlated information to many observers, who are in turn observed by others. These others may neglect the induced correlation in what they observe. In consequence, correlation neglect increases the influence of agents who are more heavily connected in the social network (see DeMarzo, Vayanos and Zwiebel (2003) and the review of Golub and Sadler (2016)).

Correlation neglect is captured in a sequential quasi-Bayesian setting in Bohren (2016). In Bohren's model, states are equiprobable, signals are bounded, and there is a given probability that each agent is *social* (observes the actions of predecessors) or *nonsocial* (does not observe any predecessor actions). Whether an agent is social is unknown to others. Agents may either underestimate this probability (a form of correlation neglect) or overestimate it (which could be called correlation overestimation).[27]

In the case of rational agents, this is a model of fragility in cascades and social learning. When enough agents take the same action, a cascade forms, in the sense that a social agent follows the preponderant action of predecessors. However, when a nonsocial agent arrives, this agent may provide a social information shock by taking an action in opposition to the cascade. This is informative to later social agents, so that the initial cascade may be broken. In the rational case, eventually those agents who do observe predecessors make correct decisions, since the information implicit in the actions of an infinite stream of nonsocial agents is conclusive.

To understand outcomes in the imperfectly rational cases in Bohren's setting, let q be the probability that any agent observes the actions of predecessors, and let $\hat{q}$ be agents' perception of that probability. If the possible values of $\hat{q}$ are in an intermediate interval (an interval which includes q), then Bohren shows that, just as in the rational case, with probability one the social agents eventually make correct choices.

When $\hat{q}$ is below this intermediate interval, agents view past actions as often being taken independently, and therefore severely overestimate how informative these actions are about private signals. When, by chance, a strong enough preponderance of agents favors one of the available actions, agents fall into a cascade. So the preponderance of one action tends to grow over time. Since agents think this growing preponderance is coming largely from independent private signals, social agents grow increasingly confident in the correctness of the latest cascade. In the limit agents become

[27] Guarino and Jehiel (2013) offer an alternative approach to modeling imperfect understanding by agents of the relation between others' private information and their actions.

sure of either the wrong state or the correct one.

This possibility of strongly held faith in the wrong state provides an interesting contrast with the BHW cascades setting, in which there is failure of asymptotic learning but cascades are fragile. It also contrasts with a rational benchmark with continual arrival of nonsocial agents, in which asymptotically the beliefs of the social agents become arbitrarily strong, but always converge to the correct state.

When $\hat{q}$ is large (i.e., to the right of the abovementioned interval), beliefs fluctuate forever, so again there is not asymptotic learning. Even if, at some date, there were a very strong preponderance of action H, for example, agents would believe that this derives almost entirely from cascading by predecessors. This makes the system extremely fragile. When by chance (as must eventually happen) even a few nonsocial agents take the opposite action, the next social agent will no longer be in a cascade, and will sometimes choose L.

As Bohren points out, thinking that $\hat{q} < q$ can cause agents to have excessive faith in a sequence of identical actions relative to expert scientific opinion. In the social learning model of Eyster and Rabin (2010), correlation neglect takes a more extreme form—observers think that each predecessor decided independently based *only* upon that agent's private information signal.[28] In their model, state and actions are continuous. Beliefs about others are analogous to $q = 1$ and $\hat{q} = 0$ in the Bohren (2016) model. In consequence, the views of early agents are very heavily overweighted by late agents, convergence to the correct belief is blocked (even with sharing of continuous beliefs or actions), and agents become highly confident about their mistaken beliefs.

One lesson that comes from analyses of imperfect rationality and social learning is that biases that cause agents to be more aggressive in using their own signals, such as overconfidence, or such as overestimation of how heavily others have observed their predecessors, tend to promote the use of private signals. Under appropriate conditions,

[28] Hirshleifer and Teoh (2003) and Eyster, Rabin and Vayanos (2018) apply such neglect of the signal-dependent behavior of others to financial markets.

this can in turn improve social information aggregation. In contrast, persuasion bias tends to have an opposite effect, causing agents to defer too much to history.

6.2.3 Other heuristics and psychological biases

So far in this section we have discussed models that explicitly analyze the effects of psychological biases such as correlation neglect and overconfidence on social learning. Such models fully endogenize beliefs and behaviors. Another general approach is to make exogenous assumptions about the agent's mapping from observed actions and payoffs into the agent's actions. Ellison and Fudenberg (1993, 1995) provide pioneering analyses using this heuristic agent approach (see the Online Appendix, § A.1 for details). In recent years, behavioral economics has moved toward basing assumptions on evidence from human psychology, and endogenizing biased belief formation as part of decisions (e.g., Daniel, Hirshleifer and Subrahmanyam (1998), Rabin and Schrag (1999)). The model of overconfident information processing in social learning of Bernardo and Welch (2001) is an example of this.

Bohren and Hauser (2019) examine a setting that allows for a variety of types of possible psychological biases in social learning, including correlation neglect. In this model, signals are continuous (and may be unbounded). They focus on settings in which enough information arrives so that if agents were rational there would be asymptotic learning (via the arrival of either public signals or nonsocial agents). However, owing to psychological bias, asymptotic learning can fail, which can take the form of convergence to a mistaken action, permanent disagreement over action, or infinite cycling. For example, when agents overreact to private signals, and where there is a positive probability of nonsocial types, there can be infinite cycling between actions. When agents underreact, there can be fixation upon a mistaken action. Furthermore, when incorrect herds last forever, beliefs converge almost surely to the incorrect state. So consistent with Bohren (2016), and in contrast with the BHW model, longer herds become increasingly stable.

7 Costly information acquisition

People often have a choice of whether or not to acquire information. We next examine the effects on social learning of costly acquisition of either direct private information signals about the state (Section 7.1) or about predecessors' actions (Section 7.2).

In a scenario with exogenous private signals and information cascades, such as the SBM, the signals of late agents do not contribute to social knowledge, because once a cascade forms, such signals do not affect actions. When agents can acquire private signals, it is unprofitable to do so if the signal will not (or is unlikely to) affect the agent's action. So in such settings there is often a similar conclusion, that late agents mimic their predecessors.

If private signals are costless, then asymptotic learning occurs when private signals are unbounded (as noted in § 4), and may occur when the action space is continuous (as described in § 3). Relative to this scenario, a positive cost of observing private signals degrades learning. A uniform conclusion of several papers to be discussed is that even in settings with unbounded private signals or continuous action spaces, asymptotic learning does not occur if there is even a small positive cost of investigating. In practice, costs of gathering or processing information are likely to be positive. So these results suggest that asymptotic learning will not be achieved.

On the other hand, if private signals are costless, introducing a cost of *observing predecessor's actions* can *improve* social learning. In such a setting, an agent with a very informative signal realization may choose not to observe others' actions. Thus, her action conveys greater incremental information, which benefits later agents.

7.1 Costly acquisition of direct private signals about state

Costs of acquiring private information introduce another information externality of social learning. In deciding whether to buy a signal, agents do not take into account the indirect benefit that accurate decisions confer upon later observers.

This externality is illustrated by modifying the SBM of § 2 so that agents have a choice in acquiring private signals about the state. Each agent can pay some cost $c > 0$ and observe a binary signal with given precision p, or can pay nothing and observe no signal. In this setting, agents I_n, $n > 1$ will not acquire a signal, no matter how small the cost. To see this, suppose that the cost is sufficiently small that it pays for I_1 to acquire a signal. Then I_2's social belief is either p or $1 - p$, depending on whether I_1 chose H or L. Even if I_2 were to acquire a private signal, imitating I_1's action remains a weakly optimal action for I_2, regardless of the signal realization. Thus, the signal has value zero, so I_2 will not purchase it, and instead imitates I_1. Agent I_3 understands this, so by the same reasoning, it is not optimal for I_3 nor for any subsequent agent to acquire information, and all agents imitate I_1. So the social outcome impounds even less information than in the SBM.

In more general settings as well, when there are at least small costs of acquiring a private signal about the state, agents stop acquiring private signals, resulting in complete blockage in the growth of social information. So only a few individuals end up acquiring private signals. This is a version of the "Law of the Few" (see Galeotti and Goyal (2010)).

As several authors have shown, under appropriate assumptions, asymptotic learning occurs if and only if infinitely many agents have access to unbounded signals at an arbitrarily small cost. If agents incur even a small cost of acquiring information, incorrect cascades can arise and therefore the social outcome can be fragile.[29] We discuss such models of costly information acquisition below.

The main ideas of an early contribution on costly signal acquisition by Burguet and Vives (2000) can be seen in a simplified model with unbounded signals and continuous actions and states. Other things equal, an agent prefers an action as close as

[29]This assertion refers to a slightly generalized definition of cascades: acting irrespective of the value of a potential private signal owing to the fact that the agent chooses not to acquire the private signal.

possible to the value of a continuous state θ, as with Mean Squared Error preferences. Each agent's objective is to minimize the sum of the cost of the signal and the negative mean squared error.

The common initial prior on θ is a normal distribution with mean μ_0 and precision ρ_0. Agent I_n takes action a_n after observing the actions of predecessors and a conditionally independent signal s_n which is normally distributed with mean θ and precision r_n. Suppressing the agent subscript, each agent chooses the precision $r \geq 0$ of her private signal at cost $c(r)$, where $c(r)$ is convex and increasing.

An agent's unique optimal level of precision is easily inferred by her successors. Given this, and as actions are continuous, each agent's action perfectly reveals the agent's private signal. So the social information available to I_n is the realization of the $n-1$ conditionally independent normal signals of I_n's predecessors. Consequently, the social belief of I_n is summarized as a normal random variable θ_{n-1} with mean μ_{n-1} and precision ρ_{n-1}. Owing to normality, ρ_{n-1} is equal to sum of the initial ρ_0 and the precisions of signals of predecessors.

Burguet and Vives observe that asymptotic learning is equivalent to the requirement that social precision ρ_{n-1} increases without bound of as n increases. But as ρ_{n-1} increases without bound, I_n's marginal benefit of acquiring additional precision goes to zero. Consequently, asymptotic learning occurs if and only if $c'(0) = 0$, i.e., the marginal cost of additional precision is zero at $r = 0$. For instance, if the smallest available precision is $r_0 > 0$ at cost $c_0 > 0$, so that $c'(0)$ is infinite, then learning is incomplete. Essentially, public belief becomes exceedingly accurate due to information acquisition by a large number of predecessors. So later individuals find it unprofitable to pay c_0 (or more) to acquire a private signal. Thus, in a setting with unbounded private signals and a continuous set of actions, asymptotic learning occurs if and only if agents can acquire signals, no matter how noisy, at a cost arbitrarily close to zero.

In Mueller-Frank and Pai (2016), agents acquire information about finite samples of predecessors' actions and payoffs, and act in sequence accordingly. Agents differ in the realizations of their costs of sampling, and in their sample outcomes, both of which are private information. Asymptotic learning occurs (i.e., probability of taking the best action goes to 1) if and only if sampling costs are not bounded away from zero in the sense that costs can be arbitrarily close to zero for an unlimited number of agents.

As discussed in § 3, Ali (2018*a*) introduces a notion of responsiveness which, loosely speaking, requires that any change in an agent's beliefs changes the optimal action. Ali shows that even with responsiveness, there may not be asymptotic learning if information is costly to acquire, since the benefit of greater accuracy may not outweigh the cost of information. Responsiveness implies asymptotic learning if and only if the minimum across agents of the costs of gathering information are arbitrarily close to zero.

The consistent message from Burguet and Vives (2000), Mueller-Frank and Pai (2016), and Ali (2018*a*) is that asymptotic learning is not robust to introducing costs of acquiring (or processing) private information. Instead of costs of acquiring private information, we can consider costs of acquiring information from predecessors. Suppose that we modify the SBM so that for a small fixed cost an agent can talk to a predecessor to find out the rationale behind her action choice. In other words, the agent can learn the predecessor's belief (which may reflect information that she has acquired in conversations with her predecessors). Nevertheless, as long as there is even a small cost of such conversations, incorrect cascades occur with positive probability. Intuitively, as beliefs become increasingly informative, at some point it pays for an agent to simply follow the action of the agent's immediate predecessor rather than paying to learn the predecessor's belief. So there is not asymptotic learning.

7.2 Costly or noisy observation of past actions

It is often costly to observe the actions of others. For instance, in evaluating the decision to invest in a startup firm, a venture capitalist can devote time and effort to gathering information about the decisions of earlier potential investors.

Consider a setting in which, for a fixed cost, an agent can observe all predecessors' actions. If the cost is high enough, early agents will not incur it, and therefore will act solely on the basis of their own private signals. However, for this very reason, at some point the action history may become so informative that an agent finds it worthwhile to learn the choices of predecessors. Once this point is reached, all subsequent agents will also find observation worthwhile. So observation costs can turn early agents into sacrificial lambs, as defined in § 6.2, to the benefit of many later decision makers.

Based on this, from the viewpoint of improving the accuracy of decisions (and perhaps welfare as well), typically the observation cost should be positive but not too large. With a zero cost, as in the SBM, cascades tend to be very inaccurate. With too high a cost of observing predecessors, no agent will ever incur it, so that social learning is blocked.

This insight is developed in the model of Song (2016). Each agent I_n first observes a costless private signal and then decides whether to pay a cost c in order to observe the actions of up to $K(n) \leq n-1$ predecessors (see also Kultti and Miettinen (2007)). This generalizes the scenario just described by allowing for selective observation of predecessors. Individual decisions about which predecessors to observe build a (directed) social network of observation links. However, agents do not know the full structure of the network, as each agent's decision about which predecessors to observe is private information.

Consistent with the intuition above, social learning may improve as c increases. To see why, suppose that $K(n) = n-1$, so that each agent can observe all predecessors after paying cost c. If $c = 0$, it is optimal to observe all predecessors and results from the standard model apply: with bounded private signals, there is a chance of incorrect cascades. But if private signals are unbounded and the observation capacity is unlimited (i.e., $\lim_{n\to\infty} K(n) = \infty$), then asymptotic learning always occurs.[30]

[30]The endogenous social network has expanding observations in the sense of Acemoglu et al. (2011);

For sufficiently large c, agents who receive strong signal realizations will choose not to observe any predecessor. As in the papers summarized in § 7.1, there is no asymptotic learning for such c, because such agents decide without observing others. On the other hand, such agents increase the pool of social information, so agents who do acquire information (who will exist if c is not prohibitively large) decide correctly with probability that tends to one.

7.3 Costly Information Acquisition, Limited Observation and Groupthink

Can social observation lead to decisions that are even worse than the decisions that agents would make under informational autarky? This might seem impossible, since any information gleaned by an agent via social observation is incremental to her own private information. However, psychologists have emphasized (Janis and Mann (1977) and Janis (1982)) that "groupthink" in group deliberations causes disastrous decision failures, as if interaction with others were harming instead of improving decisions. There is also evidence suggesting that observation of others sometimes result in degradation in decision quality (a zoological example is provided by Gibson, Bradbury and Vehrencamp (1991)).

Analytically, when there are investigation costs and noisy observation of past action, agents in groups can come to decisions that are on average worse than if there were no social observation. Owing to free-riding in investigation by agents who are potentially knowledgeable, social observation can actually reduce decision quality.

To see this, first suppose that, as in the SBM, that others are observed without noise, but that there is a small cost of acquiring private signals. As discussed at the start of § 7.1, starting with I_2 all agents follow I_1, so the social belief reflects only a single signal. This is no more accurate than if agents decided independently (though welfare

see § 11.2.

is higher as agents save on investigation costs).

Suppose instead that observation of predecessors is noisy, where each agent observes binary signals about the actions of all predecessors. Suppose further that all agents observe the same binary noisy signal about the action of any given predecessor.

If the noise is sufficiently small relative to the cost of the signal, the net gain to I_2 of investigating is still negative, so she still does not investigate. But now, owing to observation noise, her action is less accurate than if she were to decide on her own. So observation of others reduces decision quality relative to informational autarky. (Nevertheless, I_2's welfare is higher than under autarky, as observation of others economizes on observation costs.)

What about later agents? Agent I_3 also just follows I_3's signal about I_1's action.[31] The same applies to all later agents, so everyone's action is less accurate than if they had decided independently. In a related setting, suppose that agents observe only the latest predecessor. In this case noise can compound repeatedly until a point is reached at which an agent again pays to acquire a private signal (Cao and Hirshleifer (1997)).

An important empirical question in social learning settings is who makes better decisions on average, the agents who follow the predominant action, or those who deviate. In the SBM (and in the general BHW cascades model), in any realization, it is the later agents who are in a cascade, and those in a cascade have observed more predecessors than those who precede the cascade. If there are many agents, then such cascading agents predominate. So in expectation, those agents who take the predominant action are better informed than deviants. This can distinguish the cascades model from other models of social influence.

[31]Agent I_3 ignores her signal about I_2's action, because she knows that I_2 imitated I_1 based on the same signal realization about I_1's action that I_3 observes. So if I_3's signal about I_2's action differs from I_3's signal about I_1's action, I_3 knows that this discrepancy *must* be caused by error in observation of I_2's action.

Notably, this prediction is reversed in the above example when there is modest observation noise and costly investigation. Now deviants are more accurate, because they acquire a signal directly, whereas cascading agents copy a garbled version of past actions. This garbles the information content of the single past action that was needed to trigger the cascade.

8 Payoff Externalities

The SBM focuses on information externalities, under which an agent's action indirectly affects others by providing them with information. Often, however, there are also payoff externalities, wherein an agent's action directly affects the payoff of another agent. We next consider the interaction between direct payoff externalities and social learning. This topic is discussed extensively in Chamley (2004*b*).[32]

We distinguish between an externality that is (i) backward looking only, or (ii) both backward and forward looking. In an externality of type (i), an agent's payoff depends only on predecessors' actions. An example is an agent's decision to join one of two queues, where the cost of waiting is increasing with the length of the queue. In such a situation, agents have no incentive to influence the inferences of later agents. Our primary focus here is on such settings.

In an externality of type (ii), an agent's payoff depends on the actions of both earlier and later agents. An example is the decision of individuals arriving in sequence to line up to get into a restaurant, if there is disutility from dining in a crowded restaurant. This can result in strategic incentives to influence subsequent agents. A literature on sequential games with learning encompasses strategic issues (see, e.g., Jackson and Kalai (1997), and Dekel, Fudenberg and Levine (2004)). Type (ii) externalities also arise in sequential voting settings in which voters who care about election outcomes are

[32]In § 9.2 we discuss pecuniary externalities, i.e., changes in the price of adoption due to predecessors' actions.

affected by the votes of their successors as well as their predecessors. We discuss this literature in §12.4.

Payoff externalities, such as network externalities or congestion, can be captured by modifying the utility function of the SBM. Let H_n be the number of agents who chose action H before agent I_n, and define L_n analogously. Consider the following utility function for I_n:

$$\begin{aligned} u(\theta, H) &= \mathbf{1}_{\theta=H} + \epsilon H_n \\ u(\theta, L) &= \mathbf{1}_{\theta=L} + \epsilon L_n. \end{aligned} \tag{4}$$

When $\epsilon > 0$, there is complementarity between the actions of different agents. This tends to reinforce cascades and herding. Recall that in the SBM, if I_2 sees a private signal contrary to I_1's action, I_2 is indifferent between the two actions and breaks the tie by choosing the action in accord with her signal. In contrast, in the setting here, and when $\epsilon > 0$, I_2 strictly prefers to imitate I_1 regardless of her private signal realization. Thus, all agents follow the action of I_1.

If $\epsilon < 0$, there are *negative* payoff externalities, such as congestion costs. Now the interplay between social learning and negative payoff externalities is more interesting. Under social learning with negative backward-looking externalities, and if the externalities (ϵ above) are not too large in absolute value, then agents, at least for a time, imitate predecessors. Veeraraghavan and Debo (2011) model congestion as the cost of waiting in a queue with random service times. The queue length conveys favorable information about the value of the service provided. As long as the waiting cost is small relative to the difference in the length of the two queues, agents ignore their private information and join the longer queue in a cascade. But when the difference is sufficiently large, the extra waiting time from a longer queue can outweigh the favorable inference.

Eyster et al. (2014) describe how backward-looking negative externalities prevent the convergence of agents to one action and improve social learning despite the continued occurrence of incorrect information cascades. We illustrate with a variation on the SBM. As before, there is a payoff component of $+1$ deriving from taking the correct action ($a = \theta$ when the state is $\theta = L$ or H). In addition, there is a negative payoff component deriving from congestion costs. The authors use a generalization of the utility function in (4). So the utility of agent I_n from taking action a at history of actions $\mathcal{F}_{n-1}$ is

$$\begin{aligned} u(\theta, H) &= \mathbf{1}_{\theta=H} - c_H(\mathcal{F}_{n-1}) \\ u(\theta, L) &= \mathbf{1}_{\theta=L} - c_L(\mathcal{F}_{n-1}) \end{aligned} \tag{5}$$

where $c_a(\mathcal{F}_{n-1})$ is the congestion cost of taking action a at history $\mathcal{F}_{n-1}$. We focus on the case where ϵ is small, which provides insight into the robustness of the SBM to the introduction of small negative externalities.

Consider two illustrative cases:

(i) *Absolute congestion costs:* $c_H(\mathcal{F}_{n-1}) := \epsilon H_n$, where H_n is the number of predecessors who have taken action H at history $\mathcal{F}_{n-1}$. Similarly, $c_L(\mathcal{F}_{n-1}) := \epsilon L_n$. Absolute congestion costs increase without bound. This is the case considered in eq. (4).

(ii) *Proportional congestion costs:* $c_H(\mathcal{F}_{n-1}) := \epsilon H_n/(n-1) \leq 1$, where H_n is the fraction of predecessors who have taken that action H at history $\mathcal{F}_{n-1}$. Proportional congestion costs are bounded above.

Under absolute but not under proportional congestion costs, costs can grow arbitrarily large. Loosely speaking, proportional congestion costs describe applications in which queues are gradually processed rather than being allowed to grow arbitrarily long. When there are proportional congestion costs of modest magnitude, on the whole the main conclusions of the SBM about incorrect cascades and herding carry through, because with ϵ small, the informational incentive to imitate outweighs opposing payoff interaction effect.

Let $d_n = H_n - L_n$ be the difference between the number of past actions H and L. For the case of absolute congestion costs, as in the SBM, an H cascade starts the first time that $d_n = 2$. (A similar analysis applies starting with an L cascade if $d_n = -2$ is reached first.) However, this cascade is temporary, as the congestion cost of action H increases over time. Eventually, an agent is reached who finds it optimal to switch to L if and only if the agent sees the signal ℓ. This agent's action reveals her signal, and an interlude of informative actions continues until another cascade starts, this time at some threshold $|d_n| > 2$.

Ultimately, permanent cascading, meaning a situation in which all agents starting from agent I_n make choices independently of their private signals, must start. Remarkably, at this point agents alternate between actions! The intuition rests upon two observations. First, as the social belief approaches 0 or 1, an agent's belief about the true state becomes less sensitive to the agent's private signal. Second, even though agents become almost certain about a state, they do not herd upon the action corresponding to that state.

To see the second point, suppose that $\epsilon = 0.1$ in item (i) above. Then even if I_n is almost certain that H is the true state, she prefers action L if $d_n > 10$, prefers action H if $d_n < 10$, and is indifferent between the two actions if $d_n = 10$. It follows that if the social belief of I_n is that the state is very likely H, I_n prefers H even after seeing signal ℓ if $d_n \leq 9$, and prefers L even after seeing signal h if $d_n \geq 10$. So agents cascade in alternation between actions H and L as d_n alternates between 10 and 9 forever. Such cascading starts with probability one, so there is never herding upon a single action.

Also, for any given ϵ there is no asymptotic learning, as in the SBM. It follows that actions are potentially fragile, i.e., sensitive to the introduction of small informational shocks. However, Eyster et al. (2014) show that when absolute congestion costs ϵ approaches zero, the system becomes arbitrarily close to achieving asymptotic learning. Of course, in applied settings the magnitude of congestion costs are often non-negligible.

What can we conclude about how congestion externalities affect social learning and information cascades? With bounded congestion costs (as in the case of proportional costs discussed above), the insights of the SBM carry through. The informational pressure to imitate eventually outweighs congestion costs, resulting in cascades and idiosyncratic behavior. In contrast, in the case of unbounded congestion costs, as occur in the case of congestion costs that are proportional to the number of adopters of an action, learning outcomes differs qualitatively from the SBM. While cascades (sometimes incorrect) occur with probability one, herding does not. Instead, agents alternate between the two actions owing to congestion costs. Once cascades start all learning ceases, so there is no asymptotic learning.

9 Social Learning in Markets

In a market for a product or financial asset of uncertain value, the decision to buy depends on the price, the agent's (buyer's) private information signal, and the decisions of predecessors that the agent has observed. This raises several questions. Does the price setting process promote or prevent cascades, including incorrect ones? How should a seller manage the social learning process? How does social learning affect market efficiency? What are the welfare consequences of social learning and cascades? We first discuss the case of monopoly pricing, in which the seller chooses prices to maximize expected profits, and then turn to competitive price-setting.

9.1 Monopoly

A monopolist may have an incentive to set price low enough to induce a cascade of buying. The dynamics of prices and buying depend on whether the monopolist must commit to a single price or can adjust prices in response to observation of the purchase decisions of early potential buyers. We first discuss the fixed price case, which can also apply to products with menu costs (costs of changing prices; Sheshinski and Weiss (1977)). It also applies to the sale of equity shares of a firm in an Initial Public Offering (IPO), since a fixed price per share is mandated by U.S. law. Much of this literature focuses on the case of an uninformed seller, i.e., a seller who has no private information about the state. In § 9.2 we consider informed traders in competitive markets.

9.1.1 Fixed Price Case

As in Welch (1992), consider an uninformed risk-neutral monopolist who offers to sell one unit of a product to each agent in a sequence at a fixed price for all buyers until the monopolist's supply of the product, n units, is exhausted. The monopolist's cost of production is normalized to zero. As in the SBM, each agent receives a binary private signal about the state $\theta \in \{0, 1\}$, which is the unknown value of the product, and can observe the choices of all predecessors. The net gain to adopting (buying the product) is the difference between the state and the price. This contrasts with the SBM, in which the net value of adoption is exogenous.

The monopolist is risk neutral and does not discount the future. As in the SBM, each agent receives a binary private signal about the state, which is the unknown value or quality θ of the product, and can observe the choices of all predecessors.

We assume that when indifferent the customer buys. If the price is sufficiently low, all agents buy independent of their private signals, which is is an information cascade of buying. At a somewhat higher product price, an agent's choice depends on her private signal, in which case her choice reveals her private information to subsequent agents. If the price is high enough, a non-buying cascade occurs, but such a price is never optimal for the seller.

Consider three possible prices ($P = P_\ell, P_0$, and P_h), where:

- $P_\ell = \mathbb{E}[\theta|\ell] = 1-p$. The first agent starts a buying cascade, yielding the monopolist a per-buyer expected net revenue of P_ℓ for the first n buyers (and zero thereafter).
- $P_0 = \mathbb{E}[\theta] = \frac{1}{2}$. A buying cascade ensues when and if the difference between the number of buys and the number of sells reaches 1. A non-buying cascade ensues

when and if this difference reaches -2.

- $P_h = \mathbb{E}[\theta|h] = p$. A buying cascade starts if and when the buy/sell difference reaches +2. A non-buying cascade ensues when and if the difference reaches -1.

From the monopolist's perspective, demand is fragile. Just a few early agents with negative signals would cause buying to collapse. For a sufficiently low signal precision p, the profit-maximizing price is P_ℓ, so that all agents buy.

Intuitively, with low precision, $P_\ell = 1-p$ is only slightly below 0.5, so it is not worth risking collapse of demand for slightly higher prices P_0 or P_h. Since $P_\ell < E[\theta]$, the seller underprices the product. This implication is consistent with the empirical finding of underpricing in IPO markets (Ritter and Welch (2002)). For higher precision p, raising the price from P_ℓ to P_0 is worth the risk, so that there is not underpricing.

In a setting with a uniform prior on θ, Welch (1992) shows that when the seller also has a private signal, a seller whose private signal indicates higher quality (a high-quality seller, for short) sets a higher price (with higher failure probability) to separate from a lower-quality seller type. Welch (1992) focused on cascades in the IPO market. Empirically, Amihud, Hauser and Kirsh (2003) find that IPO opportunities for investors tend to be either heavily oversubscribed or undersubscribed, with almost no IPOs in between. This is consistent with information cascades, in which there is positive feedback from early investor decisions to later ones.

9.1.2 Flexible Price Case

A monopolist may be able to change prices after observing each agent's buying decision. In Bose et al. (2008), the seller is risk neutral, uninformed, and can modify the price after observing each buyer's decision.[33] The relevant prices for the seller to consider

[33]In a related paper, Caminal and Vives (1996) and Caminal and Vives (1999) study social learning about product quality via market share in a duopoly. Newberry (2016) studies empirically how fixed versus flexible pricing regimes affects buyer social learning and seller profits.

(in the SBM as modified above) are a low price that leads to an immediate sale and a higher price that results in a sale only for a high signal. For each buyer, the values of the low price and the high price depend on the actions of preceding buyers.

Bose et al. show that the seller starts with a price that induces the first buyer to reveal her private signal. Once enough information is revealed, the value of additional information revelation to the seller is low. Eventually, the seller fixes a low price that induces a buying cascade. As the seller's discount factor increases, the value discovery phase becomes longer, and more information is revealed. In the limit, if the seller has no time discount, there is complete value discovery, and the seller earns $\mathbb{E}[\theta]$ per buyer. This is the best conceivable asymptotic outcome for the seller, as rational buyers will never, on average, pay more than their ex ante expected valuation.

9.2 Competitive Markets

In contrast with the monopolistic case, it is not immediately clear whether cascades will occur under competitive price setting in product or securities markets. In securities markets, when an agent buys or sells based on her private information, market prices should change to reflect at least some of the agent's private information. This makes it less attractive for an observer to imitate the trade, which opposes the formation of a trading cascade.

To see the consequences of this effect, consider a setting in which each trader receives a signal about an object with value $\theta = 0$ or 1. The trader can buy one unit at the market makers' ask price A, sell one unit to the market makers at bid price $B \leq A$, or not trade. If market makers are uninformed, it must set A and B such that no trade occurs, as otherwise the market maker would lose money in expectation. This follows from standard no-trade results for securities markets with no noise traders (Glosten and Milgrom (1985)). Thus, trivially, there is a no trading cascade.

However, in settings with noise traders (Glosten and Milgrom (1985)), market makers can profit at the expense of noise traders, and bid-ask spreads are set to accommodate trading. The adjustment of the competitive market price to reflect private information discourages the occurrence of buy or sell cascades by making it optimal for traders to use their private information. To understand why, consider a hypothetical ask price that causes a cascade in which the informed agent buys even if she had an adverse signal. This cascading would cause market makers on average to lose money, inconsistent with equilibrium.[34]

Avery and Zemsky (1998) illustrate this point in a simple setting with noise traders:

- Each trader (agent) I_n trades only once, at date n, taking one of three actions: buy one share, sell one share, or hold (do not trade).
- Traders buy from or sell to perfectly competitive risk-neutral market makers. Consequently, bid and ask prices are set so that market makers break even.

The value of the asset is $\theta \in \{0, 1\}$. With probability $1 - \mu$, trader I_n has private information (with signal realizations $s_n = \ell$ or h, as in the SBM) about the value of the asset and with probability μ, I_n is a noise trade, where a noise trader is a mechanistic agent who buys, sells, or holds with probability 1/3 each. Thus there are two types of traders, noise traders and informed traders. Traders' types are independently distributed and privately known.

There are two prices in each period n: an ask price A_n at which I_n may buy the stock, and a bid price B_n at which I_n may sell the stock. An informed I_n's utility

[34]Market makers lose money because in such a cascade an informed agent always buys, making trades uninformative. A competitive market maker sets the bid equal to the ask, since the market maker is compelled to set the price equal to the conditional expected value of the fundamental given the order. With no spread, the market maker makes no money trading with noise traders. Furthermore, the market maker on average loses money to the informed agent. (The informed agent with an adverse signal must break even when buying, so an informed agent with a favorable signal will strictly profit from buying.) It follows that the market maker would not participate.

from buying is $\theta - A_n$, and from selling is $B_n - \theta$. Let $\mathcal{F}_n$ be the publicly observed history of trades by agents $I_1, I_2, \ldots, I_{n-1}$. An informed I_n sells if $\mathbb{E}[\theta|\mathcal{F}_n, s_n] < B_n$, buys if $\mathbb{E}[\theta|\mathcal{F}_n, s_n] > A_n$ and holds otherwise. Perfect competition among market makers implies that the bid and ask prices satisfy

$$B_n = \mathbb{E}[\theta|\mathcal{F}_n, a_n = \text{Sell}] \ \leq \ \mathbb{E}[\theta|\mathcal{F}_n] \ \leq \ A_n = \mathbb{E}[\theta|\mathcal{F}_n, a_n = \text{Buy}], \tag{6}$$

where a_n is I_n's action.

In an information cascade, the action of an informed trader is uninformative as it does not depend on her private information. A noise trader's action is always uninformative. Thus, if an information cascade were to start, a competitive market maker would not be able to charge a higher ask price than the bid price. I.e., (6) would be satisfied with equality, $A_n = B_n = \mathbb{E}[\theta|\mathcal{F}_n]$. But then informed traders would sell if $s_n = \ell$ and buy if $s_n = h$ because

$$E\theta\mathcal{F}_n, \ell < B_n = \mathbb{E}[\theta|\mathcal{F}_n] = A_n < \mathbb{E}[\theta|\mathcal{F}_n, h],$$

which is a contradiction. Hence, information cascades do not form. Since there is no cascade and I_n's action has information content, the inequalities in (6) are strict. Moreover, an informed I_n buys if $s_n = h$ and sells if $s_n = \ell$ regardless of the public history $\mathcal{F}_n$. Over time, bid and ask prices converge to the true value of the stock and volatility of the stock price decreases.

We have seen that as prices adjust based on social information, there are no cascades in this setting. Nevertheless, there is a sense in which informed investors may act in opposition to their own signals: for $n > 1$, I_n may take a different action than she would have taken if she had moved first and had not seen any social information. Formally, Avery and Zemsky define a concept which they call "herds" that is adapted to financial markets. An informed trader is in what we call a *momentum herd* if the trader's optimal action is contrary to the optimal action the trader would have taken had she moved first, i.e., if she had the same private signal realization, no social information, and faced the initial bid and ask prices. Momentum herd behavior differs from an information cascade in that the action depends on the signal realization. Thus, an informed I_n is in a *buy momentum herd* if:

(i) I_n would sell in period 1: $\mathbb{E}[\theta|s_n] < B_1 = \mathbb{E}[\theta|a_1 = \text{Sell}]$

(ii) I_n buys in period n: $A_n = \mathbb{E}[\theta|\mathcal{F}_n, a_n = \text{Buy}] < \mathbb{E}[\theta|\mathcal{F}_n, s_n]$.

Momentum herds require going beyond our modified SBM setting. In that setting, condition (i) implies that $s_n = \ell$ while condition (ii) implies that $s_n = h$. Thus, a buy momentum herd is impossible in this example. A similar argument rules out a sell momentum herd in this modified SBM.

Nevertheless, momentum herds are possible in slightly generalized settings. Avery and Zemsky present an example with three states in which momentum herds are possible. In their example the signals are not monotone.[35] In a more general treatment, Park and Sabourian (2011) provide necessary and sufficient conditions for momentum herds, and show that momentum herds are possible with monotone signals.

Cipriani and Guarino (2014) provide and test a model of momentum herds using stock market data. They modify the Avery-Zemsky model by dividing time into days, where each day consists of a finite number of trading periods. There is an asset whose fundamental value remains fixed during the trading day, and receives an independent shock at the end of each day. Agents are exogenously ordered and act once. They observe the history of prices and actions, and in addition each receives a private signal regarding the value on the day at which they trade.

As in Avery and Zemsky, information cascades do not occur (as prices adjust to reflect prior trades) and momentum herds do occur. In Cipriani and Guarino's setting we can assess the prevalence of momentum herds when the state evolves stochastically through time. They calibrate their model to NYSE stock market data and estimate that inefficiencies deriving from incorrect momentum herds constitute 4% of asset value.

[35] In this example, the posterior distribution of state conditional upon signal is not ordered by first-order stochastic dominance.

Recall that in Avery and Zemsky's setting, there are competitive market makers with noise traders, trading occurs, but there are no cascades. In contrast, if transactions costs are incurred by either the market maker or traders in a setting with competitive markets and noise traders, there can be cascades of no-trade (see Romano (2007) and Cipriani and Guarino (2008*a*)). Information asymmetry about the asset value decreases as successive traders buy or sell. Ultimately, a no-trade cascade starts when the value of an informed trader's private information is less than the cost of trading induced by the transaction cost.

Lee (1998) considers a model in which each trader incurs a one-time transaction cost that enables her to trade repeatedly based upon a single private signal. This can be viewed as a setup cost, perhaps cognitive, of learning how to trade. Temporary information blockage is possible even without exogenous public information arrival. Each agent I_n enters in period n, and after entering, can buy or sell any amount of a risky asset. Owing to the transaction cost, private information can be sidelined during several periods with no trading. This quiescent interval is shattered if a later agent trades upon observing a sufficiently extreme signal. Since multiple signal values can result in the same action the equilibrium at a given date is an example of a partial cascade as discussed in § 4. When agents suddenly start to trade, there is a sudden drop or jump in price. Lee calls this phenomenon an information avalanche.

In contrast with the preceding papers, in settings in which agents have private fundamental values of assets, information cascades of buying or selling by informed traders do occur. Even though bid and ask prices adjust to reflect previous trades, informed traders with sufficiently low private value components sell and those with sufficiently high value components buy regardless of their respective signal realizations. Private values are common for illiquid assets, such as real estate and private equity. Even for liquid assets, owing to risk aversion, an agent that is endowed with substantial holdings in a firm places less marginal value on a share than an agent with no holdings. Furthermore, even under risk-neutrality, investors who value control rights can place different values on the shares of a firm.

There are several models in which private values induce cascades in asset market trading. In the model of Cipriani and Guarino (2008*b*), assets have a private value component. The authors show that information cascades, both incorrect and correct, may occur, with no asymptotic learning. A similar result is obtained by Décamps and Lovo (2006), who consider trader heterogeneity in the value of an asset deriving from differences in risk aversion and initial endowments.

Also in contrast with the no-cascades result in Avery and Zemsky's setting, Chari and Kehoe (2004) find that if agents own investment projects and have a choice as to when to trade them, then information cascades can occur. In Chari and Kehoe (2004), agents decide when to buy or sell one unit of a risky project. An agent has the option to wait, but once an agent buys or sells, she leaves the market (becomes inactive). In each period, one randomly-selected active agent receives a private signal about the value of the project. All active agents, including those who have not yet received a private signal, may buy or sell in any period in a market with bid and ask prices set competitively by market makers. Agents observe the history of buy and sell decisions, as well as prices. There is a discounting cost of waiting, but early on, uninformed and informed agents may prefer to wait to exploit the arrival of new information. As public information accumulates, the value of further information decreases with time, so a point is reached when all active but uninformed agents take a decision (either all buy from or all sell to market makers), thereby generating a possibly-incorrect cascade of buying or selling among the uninformed.[36]

10 Heterogeneous preferences

The information that a vegetarian chose a restaurant has a different meaning from the information that a meat-lover chose it. One might expect heterogeneous preferences to

[36]This is a cascade in the sense of models with a cost of acquiring information as discussed in § 7.

aid asymptotic learning by preserving action diversity. On the other hand, heterogeneous preferences can make it harder for observers to infer signals from the actions of predecessors.

If the preference types of past decision makers are common knowledge, then it is straightforward to draw an inference from a predecessor's action about her information signal. If, instead, agents have private information about their preference types, observers need to disentangle past private signals from preferences. We focus on the case of ignorance of others' preferences, under the assumption that preference types are independently distributed.

In this setting, even when signals are unbounded, there is no guarantee of asymptotic learning. Indeed, there may be no social learning at all, as is seen in the following example. There are two equally likely states, two equally likely preference types, and two actions. An action either matches or mismatches the state. The first type wants to match the state, and the second type wants to mismatch it. Each agent draws a conditionally independent signal, possibly unbounded, from the same distribution. For each signal realization, the two types take opposite actions as their preferences are opposed.

As the two preference types are equally likely, conditional upon either state, I_1 has an equal probability of choosing either action. So I_1's action is uninformative. It follows that I_2, and by similar reasoning all later agents, also have equal probability of taking the two actions, and there is no social learning.

In a model of social learning with heterogeneous preferences, Smith and Sørensen (2000) show that if agents have a finite number of preference types and signals are bounded, then asymptotic learning does not occur. This failure can take the form of information cascades, limit cascades, or confounded learning as in the example above.

In contrast, Goeree, Palfrey and Rogers (2006) find that if there is a continuum of preferences types, then asymptotic learning is possible. An agent's payoff from taking an action is the sum of a common value, which is imperfectly known, and her private value. Agents receive private signals about the common value. When the range of possible private values is greater than the range of possible common values, an agent's action is always informative, i.e., cascades do not occur, and asymptotic learning obtains.

Zhang, Liu and Chen (2015) compare a simple version of the heterogeneous preferences model of Goeree et al. (two actions, two states, and uniformly distributed private values) with the SBM (homogeneous preferences). For sufficiently large n, I_n has a higher probability of making a correct inference about the state under heterogeneous preferences than under homogeneous preferences. Zhang et al. argue that the homogeneous preferences case may fit applications such as a social media network of friends, whereas the heterogeneous preference case applies to social observation among strangers. Consequently, Zhang et al. suggest that sellers of low quality products may prefer to advertise on social media networks of friends while sellers of high quality products may prefer to advertise on social media networks of strangers.

So far, we have considered settings in which the true distribution of preference types is common knowledge. More generally, Frick, Iijima and Ishii (2020) find that misestimation of this distribution (psychological bias) severely hinders social learning. In their model, a large population of agents chooses binary actions repeatedly in each discrete period. Action payoffs depend on the continuous state and on the agent's type. Initially, each agent observes a single private signal, which may be unbounded. In each subsequent period, agents are randomly selected to meet in pairs, with each observing the action that was taken by one other agent in the preceding period.

As a benchmark case for this setting, if agents correctly understand the type distribution, then there is asymptotic learning. In the spirit of Goeree, Palfrey and Rogers (2006), random heterogeneous preferences preserve action diversity. This allows the information content of agent private signals to be revealed over time.

However, even arbitrarily small amounts of misperception break asymptotic learning. In the long run, agents approach full confidence in one state, regardless of the actual state. Intuitively, when the action space is continuous, small misperceptions can repeatedly compound with successive random drawings and observations, so that misperceptions ultimately induce extreme beliefs. This contrasts with Bohren and Hauser (2019) (discussed in § 6.2.2), who find failures of asymptotic learning only when psychological bias is sufficiently strong. It will be useful for future research to delineate more fully the circumstances under which small bias iterates to eventually generate very large effects on social outcomes.

As discussed in § 7, in models with homogeneous preferences, even with responsiveness (as defined in § 3), asymptotic learning occurs only if costs of gathering information are arbitrarily close to zero across agents. One might think that when preferences are heterogeneous, incorrect cascades would tend to be occasionally dislodged by the arrival of an agent with deviant tastes. This suggests that even without responsiveness, if information costs are small, there may be asymptotic learning. However, this need not be the case. In a model with heterogeneous preferences and a cost of acquiring private information, Hendricks, Sorensen and Wiseman (2012) find that learning can be incomplete.[37] In Bobkova and Mass (2020), each agent can acquire two pieces of costly information: (i) about a common value and (ii) about the agent's private value. Her payoff is the sum of the two values. Once the precision of the social information about the common value exceeds a threshold (due to information acquisition by earlier agents), later agents invest only in the acquisition of information about their private values. This precludes complete learning about the common value.

11 Cascades on social networks

Social networks—from word of mouth networks in iron-age villages to modern online social media websites—play an important role in the spread and aggregation of infor-

[37] Hendricks et al. test the predictions of their model with online music market data generated by Salganik, Dodds and Watts (2006).

mation in human society. In general, what an agent learns by observing others depends on the agent's position in the social observation network. And the overall structure of the network affects aggregate social outcomes, such as whether there is asymptotic learning. So network structure is a source of empirical implications for behavior and outcomes. Accordingly, a large recent literature models social learning in networks.

Networks increase the complexity of the inferences that agents need to make. In particular, when agents do not observe the whole action history, they potentially learn about the actions of those that they do not observe from the actions of those that they do observe. Thus, calculating expected utilities may require taking into account the structure of the entire social network (Mossel, Sly and Tamuz (2015)). As this places heavy computational demands on agents, the rationality assumption becomes less plausible.

So, for tractability network economists often make strong assumptions about the geometry of the network, and for both tractability and realism, often focus on non-Bayesian agents. Nevertheless, models with rational agents provide valuable benchmarks for evaluating the effects of different heuristic behaviors or beliefs.

A key question is how the geometry of the social network affects learning outcomes. As we will discuss, a general lesson from both rational and boundedly-rational models is that egalitarian networks—loosely speaking, networks in which no agent is much more important than others in the geometry of the network—tend to facilitate social learning (Bala and Goyal (1998), Golub and Jackson (2010), Acemoglu et al. (2011), and Mossel, Sly and Tamuz (2015)).

We next discuss the spectrum of models of social learning on networks. In § 11.2 we consider network models of rational social learning with sequential actions. In § 11.3 we study models with repeated actions, featuring rational agents in § 11.3.1 and heuristic agents in § 11.3.2.

11.1 Model Spectrum

Models of social learning on networks vary across several dimensions. We outline here the model spectrum, and in later subsections discuss the insights provided by these models.

11.1.1 Rationality

Substantial literatures examine settings with either Bayesian (i.e., rational) agents, quasi-Bayesian agents, and agents who use (non-Bayesian) heuristics.[38] We say that agents are *quasi-Bayesian* if they use Bayes' rule to update beliefs, perhaps with incorrect inputs. For example, agents may ignore some of their signals as in Bala and Goyal (1998)).

Heuristic agents are those whose action choices are far from following any expected utility maximizing decision process or whose beliefs are not broadly compatible with Bayes' rule. For example, in the early and influential model of DeGroot (1974), agents repeatedly update their beliefs to equal the average of their social network neighbors' previous period beliefs.

In settings with repeated moves, forward-looking agents may wish to take actions that in the short run yield lower utility, in order to influence their peers to reveal more information in the future (see, e.g., Mossel, Sly and Tamuz (2015)). Alternatively, there are models with myopic agents, who maximize expected utility in each period, but completely discount future utility. In consequence, they do not take into account the effects that their actions have on the actions of others.

The myopia assumption is often made for tractability, to facilitate the study of otherwise-rational agents. Moreover, the complexity of forward-looking inference is

[38] Bayesian: Parikh and Krasucki (1990), Acemoglu et al. (2011), Mossel, Sly and Tamuz (2015) and Arieli and Mueller-Frank (2019). Quasi-Bayesian: Bala and Goyal (1998), Molavi, Tahbaz-Salehi and Jadbabaie (2018). Heuristic: Golub and Jackson (2010).

a reason for why in practice agents may actually behave myopically. In some papers, agents are effectively made myopic by assuming that an agent's action has no effect on the behavior of others.[39]

11.1.2 Sequential Single Actions vs. Repeated Actions

Early social learning models assumed that each agent acts only once in an exogenously determined order (Banerjee (1992) and BHW). Most of the network literature follows suit (notably Acemoglu et al. (2011)); other models consider agents who act repeatedly. A pioneering example of a quasi-Bayesian network model with repeated actions is that of Bala and Goyal (1998). An early repeated action model with rational myopic agents on a social network is that of Parikh and Krasucki (1990), who generalize the two agent model of Geanakoplos and Polemarchakis (1982) to a network setting. Later repeated action myopic models include Rosenberg, Solan and Vieille (2009) and Mossel, Sly and Tamuz (2014). Forward-looking agents acting repeatedly on social networks were studied by Mossel, Sly and Tamuz (2015).

11.1.3 Signal Structure, Action Space, and State Space

We next consider other assumptions about signal structure, the action space, and the state space. Some models allow for unbounded signals or non-atomic signals, as in Smith and Sørensen (2000). Unless otherwise mentioned, the results we discuss in this section apply to general signals, under the assumptions of a binary state and binary actions as in the SBM.

[39] Papers based on the myopia assumption include Parikh and Krasucki (1990), Mossel, Sly and Tamuz (2014), and Harel et al. (2021). In Gale and Kariv (2003*a*), agents act myopically because there is a continuum of agents in each node of the network.

11.1.4 Information about Whom Others Observe

In the basics cascades setting and many other social learning models, there is common knowledge that everyone observes the full action history. Even in settings where agents do not observe the full action history, it is often assumed that each agent knows exactly who her predecessors have observed. More generally, an agent's neighborhood of observation can be private information of that agent. In the model of Acemoglu et al. (2011), agents know the distribution from which the neighborhoods of other agents are drawn, but not the realizations. In the imperfectly rational model of Bohren (2016) discussed in § 6.2, there is a chance that any agent has an empty neighborhood. Misestimation of this probability by others leads to failures of asymptotic learning.

11.2 Sequential Actions

In this section, we consider models of rational social learning with sequential actions on networks. Banerjee (1992) and BHW assume that every agent observes the actions of all predecessors in the queue. This is a simple network structure wherein agents can be identified in order of moves with the positive integers, and each agent I_n observes the actions of all of her predecessors I_m, where $m < n$.

A subsequent literature retains the exogenous ordering of actions, but relaxes the complete observation structure, so that each agent observes only a subset of her predecessors. We discussed some models with this feature in § 6.

To further consider such settings, let I_n's *neighborhood*, N_n, be the set of agents whose actions agent I_n observes before acting. Çelen and Kariv (2004*b*) study a model in which $N_n = \{I_{n-1}\}$: each agent observes her immediate predecessor. In their model, the state is equal to the sum of the agents' private signals, rather than being binary as in the SBM. With this state and network structure, neither herding nor information cascades arises, but the probability that later agents mimic their immediate predecessors tends to one.

Acemoglu et al. (2011) introduce a general network structure: the neighborhood N_n of agent I_n can be any subset of $\{I_1, \ldots, I_{n-1}\}$, and, moreover, can be chosen at random, exogenously and independently. They study asymptotic learning (as defined in § 4.3): under what conditions does the probability that agent I_n takes the correct action tend to 1 as n becomes large? Part of the answer is that asymptotic learning never occurs when agents can observe all predecessors and have bounded signals (BHW, Smith and Sørensen (2000)), owing to incorrect information cascades (or limit cascades). Nevertheless, Acemoglu et al. (2011) show that asymptotic learning is possible even with bounded signals in some networks with incomplete observation structures.

For example, the presence of sacrificial lambs as defined in § 7—agents who are unable to observe others—can induce asymptotic learning. To see this, suppose that $N_n = \{\}$ with probability $1/n$, and that $N_n = \{I_1, \ldots, I_{n-1}\}$ with the remaining probability $(n-1)/n$. Sacrificial lambs act according to their private signals only. The rest observe all their predecessors. The sacrificial lambs choose the wrong action with a constant probability that does not tend to zero with n. But these mistaken actions become exceedingly rare, as the frequency of sacrificial lambs tends to zero. Furthermore, their actions reveal independent pieces of information to their successors. Because the probability of arrival of a lamb decays slowly enough, there are infinitely many sacrificial lambs, and the rest eventually choose correctly with probability one.[40] Acemoglu et al. (2011) provide a more general condition that ensures asymptotic learning. As in the sacrificial lambs example, this condition applies to stochastic networks only; indeed, deterministically placed sacrificial lambs cannot produce asymptotic learning, since asymptotic learning requires that *all* late players choose correctly with high probability.

Acemoglu et al. also show, conversely, that asymptotic learning cannot hold when some agents play too important a role in the topology of the network. This happens when there is a set of important agents $\{I_1, \ldots, I_M\}$ that constitute the only

[40]The beneficial effect of sacrificial lambs is similar to an effect in Bernardo and Welch (2001) and Bohren (2016) as discussed in § 6.2.3, wherein imperfectly rational agents act based on their private signals instead of cascading.

social information for an infinite group of agents. When their signals happen to indicate the wrong action—which occurs with some positive probability—infinitely many agents follow suit. Along these lines, Acemoglu et al. say that a network has *non-expanding observations* if there is some M and $\varepsilon > 0$ such that, for infinitely many agents I_n, the probability that N_n is contained in $\{I_1, \ldots, I_M\}$ is at least ε. In this case, asymptotic learning does not occur. The lesson that important agents impede the aggregation of information is one that—as we shall see—recurs frequently across a wide spectrum of models.

Turning to short-term dynamics, Acemoglu et al. (2011) show that social learning can sometimes induce beliefs that are contrarian with respect to a subset of predecessors, and anti-imitation (see also Eyster and Rabin (2014) discussed in § 6.2). Intuitively, suppose that both I_3 and I_2 observe I_1 only, and that I_4 observes I_1, I_2 and I_3. Then I_4 should place positive weight on I_3 and I_2, and negative weight on I_1 to offset double-counting. So, within a broad stream of imitation, there can be eddies of contrarian behavior.[41]

Overall, these papers show that sometimes smart observers do not just follow the herd. Sometimes, smart agents may put much greater weight on the actions of fewer agents (resulting in following the minority, as in Callander and Hörner (2009), discussed in § 6.1), or even put negative weight on some (as in Acemoglu et al.). Although such effects are far from universal properties of social learning models, such findings provide a useful caveat to the intuitive notion that agents put positive and similar weights on the actions of predecessors.

Until now we have considered network models in which the neighborhoods N_n are either deterministic, or drawn independently. Less is known about the case in which neighborhoods N_n are not independent (as analyzed by Lobel and Sadler (2015)). For

[41]For a formal example, see "Nonmonotonicity of Social Beliefs" in Acemoglu et al. (2011, Appendix B). Such contrarian behavior does not occur in networks in which each agent can only observe one predecessor (a *tree network*).

example, a given agent may have a chance of being observed by either everyone or by no one.

In Arieli and Mueller-Frank (2019), agents are placed on a two (or higher) dimensional grid, and the timing of their actions is given by their distance from the origin. The observation structure is chosen at random according to a parameter p; each agent is independently, with probability p, *connected* to each of her grid neighbors who are closer to the origin. An agent I_n observes an agent I_m if there is a path of connected agents starting from I_n and ending in I_m. This is a special case of Lobel and Sadler (2015), but not of the Acemoglu et al. (2011) setting, since the realized neighborhoods are not independent. Specifically, if I_n is not observed by any neighbor, then she would not be observed by any other agent, and so these events cannot be independent.

After the observation structure is realized, the agents take actions sequentially, according to their distance from the origin. Arieli and Mueller-Frank (2019) study what they call *α-proportional learning*, meaning roughly that at least an α fraction of the (infinitely many) agents choose the correct action. More accurately, this obtains whenever an α-fraction or more of the agents in the ball of radius r choose the correct action with probability that tends to one as r tends to infinity.

Fixing $\alpha < 1$, Arieli and Mueller-Frank (2019) ask for which values of p is α-proportional learning obtained. When $p = 1$, each agent observes all agents who lie between that agent and the origin. In this case cascades form as in the SBM, so α-proportional learning does not hold. Their main result is that, nevertheless, for all sufficiently large $p < 1$, α-proportional learning *does* hold.

As in Acemoglu et al. (2011), this conclusion is based upon sacrificial lambs. For any $p < 1$ there is a constant fraction of agents who observe no other action. The actions of these agents provide independent information to observers. As $p \to 1$ these agents become more rare, but the network becomes more connected, delivering this information to a larger and larger fraction of the population. Thus, learning is achieved as long as there are some sacrificial lambs, regardless of how small their fraction is.

An interesting open question is whether there exists a *deterministic* network structure in which asymptotic learning is attained for some bounded private signal distribution. The sacrificial lamb mechanism by which asymptotic learning is achieved in *stochastic* networks is simple, but does not seem to have an obvious analogue in deterministic networks.[42] Thus, learning in deterministic networks would have to result from a different mechanism; it is interesting to understand whether such mechanisms exist.[43]

11.3 Repeated actions

Often in practice people take actions more than once while observing and learning from other bidirectionally. For example, in online social networks people observe product choices and lifestyle choices repeatedly over time. To understand such interactions, we next study models in which agents act repeatedly. As a major deviation from the sequential learning setting, models of repeated action require different techniques and generate new insights. We first examine rational settings, and then turn to imperfect rationality.

11.3.1 Rational models

The logic of information cascades can still sometimes apply in settings with repeated actions, such as when observation structure is unidirectional (contains no loops): i.e., each agent I_n observes only actions of earlier-arriving agents I_m, where $m < n$. This can be seen simply by generalizing the SBM to allow each agent I_n to take actions every

[42] In a deterministic network, if there is an infinite number of sacrificial lambs, then the probability that I_n makes the correct choice cannot tend to one, since the sacrificial lambs have a fixed positive probability of taking the wrong action.

[43] An alternative possible definition of asymptotic learning is that the probability that the majority of the first K individuals takes the correct action approaches one. Under this definition, deterministic sacrificial lambs can induce asymptotic learning if they are sparse enough, e.g., at positions 1, 10, 100 and so on.

period starting from date n, but maintaining the unidirectional observation structure. In this setting, actions choices are identical to choices in the SBM, because after date n, I_n has nothing new to learn from earlier arrivers, and never observes later arrivers.[44] So there are still incorrect information cascades and a failure of asymptotic learning.

This analytically trivial extension of the SBM illustrates that inefficient information aggregation in the SBM does not rely upon the fact that agents move only once. Instead, the problem lies in the observation structure. There is an undue observational focus on a limited set of agents—early ones. This suggests that in more general settings with repeated moves—even those with observation loops—having a bad observation structure, such as one that is far from egalitarian, may produce adverse outcomes.

Under highly connected observation structures (such as one where there is a path of observation between any pair of agents), the concepts of herd behavior and information cascades need to be generalized to apply formally. Following the literature, we consider more general notions of agreement that capture phenomena that are similar to herding. Here *agreement* is the situation in which all agents eventually agree about what action is optimal (though they need not have identical beliefs). Likewise, we consider notions of *learning* (or the failure thereof) that capture repeated action analogues of information cascades. Indeed, herding in sequential models is a form of agreement, since it occurs precisely when all agents (except a finite number, out of infinitely many) agree on the optimal action. In an analysis covering a wide class of social learning models, Mossel et al. (2020) find that agreement and asymptotic learning are closely related, and provide conditions under which each of these phenomena occur.

An incorrect information cascade implies that information is lost and that agents do not learn the true state. In repeated action models—in lieu of incorrect information cascades—we ask whether agents learn the state, and we study the mechanisms that

[44] Speaking more carefully, this is an equilibrium because if agents I_1 through I_n keep repeating the same actions forever after each one arrives, then I_{n+1} learns nothing new when they repeat their actions, and hence I_{n+1} also repeats the same action forever.

drive failures to do so. These mechanisms often resemble cascades in that some private information is permanently lost, even though agents do not completely disregard their own signals in the early periods.

Information is lost for a reason that is similar to models of cascades in settings in which each agent acts only once. A point is reached at which an agent's signal is incrementally informative relative to social information, but owing to coarseness of the action space, the agent's action does not adjust to reflect this incremental information. So when social information is sufficiently precise, actions do not reveal all the private information to others.

The study of agreement in social learning goes back at least to Aumann (1976). This seminal paper showed that two agents who have a common prior, who receive a signal regarding a binary state, and who have common knowledge of their posteriors, must have equal posteriors; agents cannot "agree to disagree."

Aumann's model does not consider the dynamics of how agents arrive at common knowledge of posteriors. Several subsequent models study how agreement is reached via social learning. Geanakoplos and Polemarchakis (1982) consider two agents who observe private signals regarding a state, and then repeatedly tell each other their posteriors. The authors show that the agents reach common knowledge of posteriors, and hence their posteriors will be identical.

Parikh and Krasucki (1990) extend this conclusion to a network setting. They consider a finite set of agents connected by a network. In each period, each agent I_n learns the posteriors of the members of her neighborhood N_n. Under the assumption that the network is *strongly connected*,[45] Parikh and Krasucki show that the agents reach common knowledge of posteriors, and hence agree.

This result was extended by Gale and Kariv (2003*a*) to a setting in which each

[45]The network is strongly connected if there is a chain of neighbors connecting every pair of agents: formally, for each two agents I_n, I_m there is a sequence of agents $I_{n_1}, I_{n_2}, \ldots, I_{n_{k+1}}$ such that $I_n = I_{n_1}$, $I_m = I_{n_k}$, and $I_{n_{\ell+1}} \in N_{n_\ell}$ for $\ell = 1, \ldots, k$.

agent receives a signal at the initial date and then takes an action in each of infinitely many periods. Agents face a common decision problem at each period: they choose an action which results in a stage utility that depends on their action and the unknown state.[46] At each period, they observe their neighbors' actions, but not the stage payoffs. The agents are assumed to be myopic, so that at each period they choose an action that maximizes their stage utility. Neighboring agents do not necessarily eventually agree on actions, even at the limit. But any disagreement is due to indifference. Specifically, suppose that I_n and I_m observe each other, I_m takes an action a infinitely often, and I_m takes an action b infinitely often. Then both must must, at the limit, be indifferent between a and b. This result was extended by Rosenberg, Solan and Vieille (2009) to forward-looking agents who maximize their discounted expected utility, rather than myopically choosing an action in each period.

The mechanism at work here is the *imitation principle* (also known as the improvement principle; see Golub and Sadler (2016)), which asserts that in the long run an agent will be able to do at least as well as an agent that she observes, i.e., a neighbor. Agent I_n is always free to imitate I_m, i.e., choose at each time period the action that I_m chose in the previous one. Since I_n myopically maximizes her expected stage utility, her expected stage utility at time t is at least as high as I_m's expected stage utility at time $t-1$. Since the network is strongly connected, it follows that at the limit $t \to \infty$ all agents have the same expected stage utilities.

Now suppose that with positive probability agent I_n thought that agent I_m's action were strictly suboptimal, at the limit $t \to \infty$. Then I_m would have lower expected utility, in contradiction to the imitation principle. So even if I_n disagrees with I_m, I_n must believe that I_m's action generates equal utility to the action that I_n chooses. That is, I_n is indifferent between these actions.

While neighboring agents that disagree must be indifferent, non-neighbors can disagree without indifference. Consider a network of four agents I_1, I_2, I_3, I_4, who

[46] A *stage utility* is a payoff received in a particular period of a dynamic game.

are connected along a chain, with each observing both the agent's predecessor and successor (if there is one). It is possible for I_1 and I_2 to converge to action L, while I_3 and I_4 converge to H. In this case I_2 and I_3 must be indifferent between L and H, but it is possible for I_1 and I_4 to not be indifferent. This happens in the case of binary signals and actions, when I_1 and I_2 receive an ℓ signal, I_3 and I_4 receive an h signal, and the agents' tie breaking rule is to stick to their previous period action. In this case the agents' actions all immediately converge. After seeing that agent I_3 does not change her action, I_2 concludes that I_4 got an h, and thus I_2 becomes indifferent. Likewise, I_3 becomes indifferent. But I_1 does not know that I_2 is indifferent: from I_2's point of view, it may well be possible that I_3 is also taking action L. Thus I_1 is not indifferent, and neither is I_4, and yet they disagree.

This raises the question of whether the possibility of indifference is robust. A partial answer is that when the action set is rich enough, such indifference is impossible, and the conclusion is again that agents must converge to the same action. Mossel, Sly and Tamuz (2014) show that even for a binary action set, when private signals induce non-atomic beliefs (i.e., no belief occurs with positive probability) then again such indifference is impossible, and hence all agents must converge to the same action.

Even when the agents do converge to the same action, it may be an incorrect one, so that the agents do not learn the state. In a quasi-Bayesian setting,[47] Bala and Goyal (1998) show that the learning outcome depends on the network geometry. An important example of a strongly connected network in which agents may not converge to the correct action is the *royal family*, in which all agents directly observe a small group, but that small group does not directly observe all others.

A similar phenomenon occurs in a setting with forward-looking, Bayesian agents who each receive a signal at period 0, and thereafter choose in each period a

[47] An important difference between the Bala and Goyal (1998) model and the others discussed here is that in their model agents have strategic experimentation incentives that determine endogenously the private signals that they observe.

binary action with the objective of matching a binary state (Mossel, Sly and Tamuz (2015)). In a network with a royal family, incorrect signals received by the royal family can cause the entire population to eventually adopt the incorrect action. When this happens, the early period actions of agents in the population are still dependent upon their own private signals, but this information does not propagate through the network. The outcome is closely related to information cascades in models with a single action, in that social information can *eventually* cause the agents to disregard their own private signals.

Conversely, Mossel, Sly and Tamuz (2015) show that in infinite networks that are egalitarian, even with bounded signals, agents all converge to the correct action. In their terms, a network is said to be *egalitarian* if there are integers d and L such that (i) each agent observes at most d others and (ii) if agent I_n observes I_m, then there is a path of length at most L from I_m back to I_n. The first condition excludes agents who obtain large amounts of social information. The second excludes royal families, who are observed by many of their "subjects" but do not reciprocate by observing their subjects directly, or even indirectly through short paths. Another interpretation of the second condition is that it requires the network to be approximately undirected, where an undirected network is one in which the second condition holds for $L = 1$.

Thus, both Bala and Goyal (1998) and Mossel, Sly and Tamuz (2015) conclude that networks in which a subset of agents plays too important a role can hamper the flow of information and lead to failures of learning, and that asymptotic learning occurs in networks in which all agents in the network play a similar role.[48] The mechanism underlying the failure of asymptotic learning in non-egalitarian networks is similar to the cause of cascades in the SBM: many agents choose actions that are heavily influenced by social information, and thus do not reveal their private signals. In the SBM, this social information comes from the early agents. In a setting with non-egalitarian networks

[48]Interestingly, for undirected networks and myopic agents, Mossel, Sly and Tamuz (2014) show that agents always converge to the correct action, for any network topology—even non-egalitarian networks in which agents can observe arbitrarily many others.

and repeated actions, it comes from agents who become socially well-informed because of their central locations in the network.

When agents choose actions from a set that is rich enough to reveal their beliefs, information is aggregated regardless of the network structure (Mueller-Frank (2014)). As a straightforward special case, when the SBM is modified to allow a rich set of actions that reveal beliefs (see § 3), information is perfectly aggregated, in the sense that agents converge to the same belief that they would hold if all private signals were to become public.

Similarly, in a network setting, DeMarzo, Vayanos and Zwiebel (2003) offer as a benchmark case a model in which agents in a connected social network are interested in a state $\theta \in \mathbb{R}$ for which all have a uniform (improper) prior. Each agent initially observes a Gaussian signal with expectation equal to θ. In each period each agent observes her neighbors' posterior expectations of θ, and updates her posterior using Bayes' Law. Each agent's posterior is Gaussian, and thus is completely characterized by the agent's posterior expectation; the variance depends on the network structure and does not depend upon signal realizations. In this model the agents converge to the same belief they would have if they were to share their private signals, so that information is perfectly aggregated.[49]

Frongillo, Schoenebeck and Tamuz (2011) and Dasaratha, Golub and Hak (2019) consider a similar setting, but where the state exogenously changes with time, agents receive a new signal at each date, and signals have heterogeneous precisions. Since the state changes, the agents cannot learn it exactly. The efficiency of information aggregation depends on the social network structure and the signal structure. In Dasaratha, Golub and Hak (2019), information is better aggregated when there is large heterogeneity in signal precisions, which helps agents filter out stale information that is entangled with information that is more relevant for the current state.

[49]The main result of DeMarzo, Vayanos and Zwiebel (2003) describes the effects of persuasion bias, wherein agents update beliefs under the mistaken premise that other agents do not observe others.

11.3.2 Heuristic models

A highly interdisciplinary literature studies heuristic models of social learning with repeated actions on networks.[50] Indeed, since repeated action models in networks tax the rationality assumption, in such settings heuristic decision making is often the more interesting case.

In mathematics and physics, models such as the "voter model" (Holley and Liggett (1975)) and the "Deffuant model" (Deffuant et al. (2000)) use tools from statistical mechanics to study interacting agents that follow rules that are simple enough that the agents resemble interacting particles. We refer the reader to other surveys and books that cover this field, such as Castellano, Fortunato and Loreto (2009) or Boccaletti et al. (2006). In sociology, a quantitative literature studies opinion exchanges on networks, partially motivated by the question of how to measure network centrality (see, e.g., Katz (1953) and Bonacich (1987)).

The statistical mechanics approach has also been applied within economics; see, e.g., Kandori, Mailath and Rob (1993), Ellison (1993), Blume (1993), Horst and Scheinkman (2006) and the survey Durlauf and Ioannides (2010). These papers focus on how play arrives at an equilibrium, typically in a coordination game, rather than on learning by players about an unknown state.

In a well-known heuristic model of learning about an unknown state, DeGroot (1974) considers a finite number of agents in a social network. Each agent is endowed with an initial subjective prior regarding a finite state. In each period each agent reveals her belief to her network neighbors, and then updates her own belief to a weighted sum of her neighbors' beliefs. These weights p_{nm} are exogenous and do not change with time.

DeGroot finds that under mild conditions this process converges, and that all agents converge to the same belief. This limit belief is equal to a weighted average of

[50] We considered heuristic models under sequential (non-repeated) actions in § 6.2.

the initial beliefs; the weights in this average are given by the first eigenvector of the matrix p_{nm}.

Golub and Jackson (2010) use DeGroot's framework to study the aggregation of information. In their model, the state is any $\mu \in [0, 1]$. The agents each start with a prior chosen from some distribution with expectation μ, and follow DeGroot's update rule. Each agent uses equal weights for all the agents in her neighborhood. Golub and Jackson show that over time the agents' beliefs will converge arbitrarily close to μ (for sufficiently large networks) as long as no agent has too high a degree. The relevant condition is that the ratio between the maximum degree and the sum of degrees be sufficiently small, which can be viewed as another notion of egalitarianism.

An overarching conclusion from § 11 is that egalitarianism in network structure, formalized in various ways, promotes information aggregation and welfare. This lesson holds across a variety of Bayesian, quasi-Bayesian and heuristic settings. Exploring this issue more deeply is an interesting direction for future research.

12 Applications and Extensions

So far we have focused mainly on variations on the basic cascades and social learning setting that are general in scope, or involve changing relatively few assumptions. We now turn to several extensions that are tailored to specific applications. In these applications, cascades and other social learning approaches offer new insights into social outcomes in both market and non-market settings.

12.1 Team Decisions, Optimal Contracting and Organizational Design

Teams often need to make joint decisions that can benefit from aggregating the information of team members who have private interests and learn by observing the actions of other members. Incentive problems influence whether agents make use of their own private signals, and therefore whether cascades hinder the efficiency of information aggregation. Such problems can be addressed through the design of a communication network and an incentive scheme for the team. We first discuss settings with rational agents.

As illustrated in the "influencers" example of § 4.4, when better informed agents decide first, information aggregation can be especially poor. As a result, anti-seniority voting systems achieve better information aggregation, when seniority is associated with precision.[51]

Khanna and Slezak (2000) study the choice of network structure—teams versus hierarchy—for a firm that seeks to aggregate the signals of multiple employees.[52] Their analysis confirms the intuition that it can be desirable to assign agents to make decisions independently, preventing information cascades, and allowing a later observer (or more generally, observers) to obtain better social information. This is similar to the discussion of sacrificial lambs in § 6.

In Khanna and Slezak (2000), agents can acquire a private binary signal about the state, which determines the profitability of a binary action on the part of the firm, such as whether or not to adopt a project. The precision of an agent's private signal is increasing with costly private effort. Ex ante identical agents make their investigation decisions before any social observation. In equilibrium they all choose the same effort

[51] However, this is not necessarily the case when agents herd for reputational reasons. As Ottaviani and Sørensen (2001) show in a variation of the Scharfstein and Stein (1990) model, anti-seniority voting systems can perform poorly.

[52] We considered more generally the effects of networks on asymptotic learning in § 11.2.

level. Each agent's action consists of a recommendation (e.g., to adopt or reject the project). The optimal compensation contract consists of a wage, a bonus if an agent's recommendation turns out to be correct ex post, and, if agents are organized into teams, a possible team bonus that is received if the team recommendation is correct.

In teams, agent recommendations are publicly announced sequentially. Owing to social observation, in equilibrium agents ex ante tend to free ride in generating private signals. Furthermore, unless the compensation scheme is heavily weighted toward individual accuracy, agents tend to cascade on the recommendations of earlier agents. The compensation scheme may optimally accommodate such cascading. These effects limit the social information forthcoming from teams to top management for choosing the firm's action.

The alternative to teams is hierarchical organization, wherein each agent reports a recommendation to the top manager without benefit of social observation. Such hierarchies reduce free-riding in information acquisition, and prevent the formation of incorrect cascades. Hierarchies therefore dominate, assuming that direct communication among agents can be cheaply suppressed.[53]

In a multi-level hierarchy, an alternative to forcing early agents to make recommendations without observing each other is to artificially coarsen their recommendations. For example, the recommendations of 3 agents can be aggregated into a single overall adopt/reject recommendation on a project. This makes it unclear to an observer whether the preponderance in favor of an action was strong (3 to 0) or weak (2 to 1). Such coarsening can cause a supervisor to whom the agents report to use the supervisor's own private information instead of cascading. This makes the supervisor's recommendation to the top layer of the firm informative. In some cases coarsening tends to improve the firm's decision (Arya, Glover and Mittendorf (2006)).

[53]In a similar spirit, Sgroi (2002) analyzes the benefits to a social planner or a firm of forcing some agents (such as a subset of customers) to make their decisions early. In Sgroi's model, promotional activity by firms encourages some customers to decide early, which can be useful for increasing sales.

12.2 Stigma, Prestige, and Related Phenomena

As mentioned in the introduction, cascades may apply to adverse inferences drawn by participants in labor markets and the market for kidney transplants, which we now discuss in turn.

Job seekers are often advised to try to avoid having gaps in their resumes, which might be indicative that previous potential employers have rejected their applications. Kübler and Weizsäcker (2003) model such unemployment stigma in a setting in which employers receive noisy signals about the quality of potential employees. The Kübler and Weizsäcker model differs from the basic cascades setting in that both employers and workers can take costly unilateral actions that modify signal precision. Kübler and Weizsäcker identify an equilibrium in which only cascades of rejection, not acceptance occur. There is indeed evidence, including field evidence, that employers view gaps in resume as indicating rejection by other employers (Oberholzer-Gee (2008), Kroft, Lange and Notowidigdo (2013)).

A mechanism very similar to stigma in labor markets can result in refusal of kidneys by patients who need kidney transplants. Patients are sorted into a queue, in order of severity of their conditions. When a kidney becomes available, patients are offered it in sequence. Each patient is offered the kidney only if all predecessors rejected it. In the model of Zhang (2010), patients have heterogeneous preferences as they differ in compatibility with an available organ and in their urgency for a transplant. There also is a common component to value, since all patients prefer a high quality kidney. Each patient's physician provides a medical judgement about the suitability of an available kidney, i.e., a private signal. If a patient is offered a kidney after refusals by earlier patients in the queue, she infers low quality of the organ, causing possible reject cascades.

Using data from the United States Renal Data System, Zhang (2010) finds strong evidence of social learning. Patients draw negative quality inferences from earlier refusals, and thus become more likely to refuse a kidney. This leads to poor kidney utilization despite a severe shortage of available organs. There is also strong evidence of social learning in other health related decisions, such as choice of health plans (Sorensen (2006)).

Research on cultural evolution has hypothesized that the human mind has evolved to confer prestige upon successful individuals (Henrich and Gil-White (2001)). In this theory, individuals who defer to an admired individual benefit from being granted access to the information possessed by that individual about how to be successful in their shared ecological environment. If different observers observe different payoff realizations, they will have different information about who is a good decision maker. So observers can acquire useful information by observing whom others defer to. The Henrich and Gil-White perspective on prestige naturally suggests that there can be information cascades in conferring prestige. If so, prestige may be a very noisy indicator of decision ability. This is an interesting topic for future research.

12.3 Social Information Use by Animals

Zoologists have developed and applied information cascades models to understand the acquisition of social information by animals for decisions about mating, navigation, predator-avoidance, foraging, and habitat selection.[54] Several empirical studies of imitation in animals conclude that social learning and cascades are important mechanisms, as distinct, e.g., from payoff externalities. In fact, imitation occurs despite the fact that payoff externalities are often negative, as imitation increases competition with others for the same food sources, territories, or mates.

Dall et al. (2005) discuss several proposed examples of animal information cascades. An extensive empirical literature documents copying in mate choices in many

[54]With regard to human habitat choice, Epstein (2010) provides a model in which there are cascades in the choice of country of emigration.

species, including humans (Waynforth (2007), Anderson and Surbey (2014), Witte, Kniel and Kureck (2015)). Gibson and Hoglund (1992) provide evidence consistent with information cascades in the choices of sage hens observing mating display groups called leks. Consistent with the logic of information cascades, there is evidence of reversal of choices upon observation of others in guppy mating (e.g., Dugatkin and Godin (1992)).

Giraldeau, Valone and Templeton (2002) apply the logic of information cascades to understand the conditions under which animals will optimally acquire information socially. The authors propose three possible examples of information cascades in nature: false alarm flights, wherein "groups of animals take fright and retreat quickly to protective cover, sometimes even in the absence of any obvious source of danger"; night roost site selection, wherein many birds adopt a site "simply because the social information indicates the site is profitable"; and mate choice copying, as just discussed. The authors argue that the possibility of incorrect cascades makes it unprofitable to expend resources to observe the actions of others.

12.4 Sequential Voting

In primary elections for nominating party candidates for U.S. presidential elections, states that go early in the process exert a disproportionate influence on the ultimate outcome. A leading explanation is that there are momentum effects wherein the beliefs of later voters were influenced by the choices of earlier voters in Iowa and New Hampshire (see Bartels (1988)).

While this finding is intuitive, in a sequential setting in which voters with private information signals seek to elect the best candidate, Dekel and Piccioni (2000) show that there exist equilibria in which voters are not influenced by predecessors' votes. In these equilibria, each voter ignores the voting history and behaves in the same way as she would in what is known as a responsive equilibrium of the simultaneous voting game. A voter is *pivotal* only if others' votes are split equally, just as in the simultaneous voting game. Since each agent's vote optimally depends on her private signal, there is asymptotic learning.

However, Ali and Kartik (2012) show that there also exist equilibria in sequential voting in which voters are influenced by predecessors, resulting in information cascades.[55] We consider a simple setting to illustrate their main ideas. Voters (agents) have identical preferences over two candidates, L and H. All voters prefer candidate θ if and only if the state is θ. The information of each voter consists of a binary private signal and the votes of her predecessors. Voters must vote for one of the candidates and cannot abstain.

Suppose that the election is decided by a simple majority vote and that the total number of voters, $2N-1$, is odd. Let $a=(a_1,a_2,\ldots,a_{2N-1})$ be the votes (actions) of all agents and $C_\theta(a)$ be the number of votes for candidate θ. In this setting, the utility of each voter is

$$u(a)=\begin{cases}1 & \text{if } C_L(a)\geq N \text{ and } \theta=L\\ 1 & \text{if } C_H(a)\geq N \text{ and } \theta=H\\ 0 & \text{otherwise.}\end{cases}$$

In sharp contrast with the SBM, in this setting every agent cares about decisions taken by others, i.e., agents are both backward-looking and forward-looking. In consequence, they care about the information their actions might convey to their successors.[56]

This raises an interesting issue: despite payoff interactions, do voters behave *sincerely*, i.e., vote for the candidate they think is best given their beliefs? Owing to the possibility of cascades, sincere voting may not be informative—i.e., agents may ignore their private signals. This raises the possibility that agents may instead want to reveal their private signals by voting in opposition to the predominant voting of predecessors.

[55] Other applications of the Ali and Kartik (2012) model include public good allocations and charitable giving.

[56] The payoff interactions here differ from those of § 8, as by assumption here there are no congestion effects nor conflicts of interest nor conformity preferences.

So it is not clear that a sincere equilibrium exists.

As Ali-Kartik show, there is indeed an equilibrium with sincere behavior. In this equilibrium, voters follow their signals unless a candidate leads by two or more votes, in which case they vote for the leading candidate regardless of their signal based on the inference that the leading candidate is probably better. This behavior results in a (possibly incorrect) cascade; there is no asymptotic learning. However, unlike in the basic cascades setting, voters care about what later voters will do, so equilibrium with sincere behavior is a strategic phenomenon.

In the cascade equilibrium in Ali and Kartik (2012), however, agents disregard the actions of predecessors who take an off-the-equilibrium-path action. As Wit (1997) and Fey (2000) note, under the reasonable assumption that a voter who breaks a cascade (i.e., off-the-equilibrium-path action) must have voted her signal, the cascade equilibrium is broken. However, plausibly the cascade equilibrium may be resurrected if there is observation error.[57]

Such an equilibrium is suggestive of evidence of voting momentum mentioned at the start of this section. This evidence indicates that when there is sequential voting, early success can promote later success owing to favorable inferences drawn by later voters. Knight and Schiff (2010) document such political momentum. They estimate that voters in early states had far greater influence than voters in later states on the outcome of the 2004 U.S. presidential primaries, and that campaign advertising choices took into account such momentum effects. Moreover, Ali et al. (2008) provide laboratory evidence that suggests possible cascading behavior in sequential voting.

Battaglini (2005) modifies Dekel-Piccione's voting model to allow for a third choice – an option to not vote. Battaglini shows that if there is even a small cost of voting, then Dekel and Piccione's result does not hold: the set of perfect equilibrium outcomes in simultaneous voting and in sequential voting are disjoint. Information

[57]Intuitively, when an agent sees a deviation from the cascade action, the observer may conclude that this is a false observation.

aggregation is worse in sequential voting than in simultaneous voting when N is large. Furthermore, when voters have a direct preference for voting for the winner in addition to a preference for voting for the right candidate, social learning can lead to a bandwagon effect (Callander (2007)).

In the papers discussed earlier, voters are motivated by the fact that they can be pivotal for the outcome, or there is a preference for voting for the winner. As other research emphasizes, other factors and motivations can also be very important, such as the cost of voting, and direct utility from voting (or not voting) as an expression of political viewpoint. Disentangling social learning from other effects is a topic for empirical research.

12.5 Changing States and Fads

In reality the best action to take often fluctuates over time, i.e., the state evolves stochastically. For example, competing web browsers sometimes leapfrog each other in functionality for users. In such a setting, the mere *possibility* that a shock to the system (that is, a change in the true value) could occur can be enough to dislodge a cascade, even if the shock does not actually arrive. For example, suppose that the state of the world is which of two restaurants has higher quality, and that this state shifts over time. The question is whether there are fads in which people shift between restaurants much more often than the underlying state switches.

To see this, suppose that just before I_{101}'s decision, with probability $1/10$ the state θ is newly redrawn from its ex ante distribution, and remains fixed thereafter. Let θ' denote the state after I_{100}, where $\theta' \neq \theta$ with probability $1/20$.

The possibility of a value shift breaks the cascade; I_{101} optimally follows her own signal. Eventually the system must settle into a new cascade—one that need not match the old one. It is easy to show that the probability of a cascade reversal is a little over $0.0935 \gg 0.05$ (see BHW). So the probability that the common behavior shifts is more than 87% higher than would be the case under full information. BHW refer to such volatile outcomes as *fads*.

What if there is a probability *each period* that the state shifts? Then if the probability is not too large, there are still cascades (including incorrect ones). However, these cascades are temporary; over time the social information in any cascade grows stale and some agent eventually returns to using her own signals (Moscarini, Ottaviani and Smith (1998)).

A fruitful direction for future research is to study what determines the prevalence of fads and the stability of conformist social outcomes with respect to external shocks that can change the underlying value state. For example, it might seem that less stable environments would be more fad-prone, but it is far from obvious that such a conjecture is true. A greater probability of state shift does not necessarily imply a greater excess probability of action shifts *relative to* the probability of state shift.

12.6 Politics and Revolutions

When citizens publicly protest or revolt against their government, the actions of early individuals convey information about the prevalence of dissatisfaction to other observers, and of the risks of government sanctions. So there can be positive feedback in submission or resistance to the regime.

In Kuran (1987), agents have both private preferences between two alternatives, and publicly declared preferences. There is a payoff/preference externality in the form of pressure toward conformity. This results in an equilibrium with high declared support for an oppressive regime that most agents privately oppose. Kuran applies this idea to the persistence of the caste system in India. Kuran (1989) emphasizes that preference falsification and positive feedback result in multiple equilibria, and that sudden change (a "prairie fire") can occur as exogenous parameters shift when a critical threshold is crossed that causes a shift from one equilibrium to another.

Lohmann (1994) explicitly models belief updating to analyze the maintenance and collapse of political regimes as information cascades. In Lohmann's model, agents have different preferred policies, and have disutility from deviation between the government's policy and the preferred ones. Agents may protest because their preferred policies may differ substantially from that of the current regime. Each agent also has a disutility component from protesting. Disutility decreases in the number of other agents who are protesting. The existing regime collapses once enough information is revealed to show that more than a critical number of individuals support an alternative regime over the current one.

Lohmann applies the model to the Leipzig protests against the communist regime in East Germany during 1989-91. In her analysis, even a small protest, when larger than expected, can reveal very strong opposition to the regime, causing the size of the protest to grow explosively. By the same token, when a large protest is expected, if the protest falls modestly short of expectations, the movement can rapidly collapse. Based on evidence from five cycles of protests, Lohmann argues that a dynamic information cascades model helps explain the successful popular uprising against the communist regime in East Germany.

Ellis and Fender (2011) combine features of Lohmann's model with those of the model of democratization, repression and regime change of Acemoglu and Robinson (2006). In Acemoglu and Robinson (2006), in order to stave off revolution, a repressive regime that is run by the rich can choose to redistribute wealth via taxation and democratize by extending the voting franchise to the poor. Some fraction of the wealth is destroyed in a revolution, which limits the incentive of the poor to rebel.

Ellis and Fender extend the Acemoglu and Robinson analysis to a setting in which the poor have private information about regime strength. In what we call the L state, a revolution would destroy a higher fraction of the wealth to be appropriated by the poor than in the H state, making it less attractive for the poor to revolt in state L. So the state of nature characterizes the strength of the regime.

In the case of interest, the rich set the tax rate such that the poor would like to revolt if they knew that the state was H and not revolt if they knew the state was L. Each poor agent (hereafter, "agent") has a conditionally independent binary private signal about the state. If the poor are not granted the franchise, agents decide in sequence whether to rebel against the regime, with each agent's decision observable to later agents. By assumption, a revolution occurs only if every agent rebels. At first, an agent with an h signal will rebel, whereas an agent with an ℓ signal will not, i.e., vetoes the revolution. However, after a sufficiently long sequence of rebels (i.e., h signals), even an agent with an ℓ signal rebels. So there are information cascades of revolution.

12.7 Legal Precedent and Information Cascades

Dating at least as far back as Oliver Wendell Holmes, a common interpretation by legal scholars of respect for precedent (*stare decisis*) is that judges acquire information from past decisions in similar cases. In the model of Talley (1999), this takes the form of imitation across courts, wherein a judge follows the decision of earlier judges "notwithstanding the facts of the case before her." The resulting loss of the private signals of later judges who respect precedent is also emphasized by Vermeule (2012, p75f). However, Talley argues that the conditions for cascades to occur in this setting are highly restrictive, in part because judges can relay information via written opinions.

Daughety and Reinganum (1999) offer a more general model of imitation across courts. In their setting, n appellate courts at the same level observe private signals about a binary state. This state indicates the single correct decision for a set of related cases to be considered by these courts. The courts make decisions in exogenous sequence. If signals are bounded, imitation across the appellate courts results in information cascades, often incorrect; even if signals are unbounded, with many courts there is a high probability of extensive imitation, resulting in poor information aggregation. The authors offer several possible examples of actual precedential cascades (pp. 161-5).

Daughety and Reinganum emphasize that incorrect cascades can be very persistent, despite the finding of BHW that cascades are fragile, because of the rarity of shocks that might dislodge a judicial cascade. As discussed in § 2, informative public disclosures can dislodge cascades. However, the authors argue that in the U.S. judicial setting, such shocks take the form of cases being brought to the Supreme Court opinions despite harmonious decisions of lower courts. Such review is rare, because the Supreme Court needs to wait for cases to be appealed, and because decisions that are harmonious across courts are rarely reviewed.

Sunstein (2009) argues that the adoption of legislation sometimes takes the form of information cascades across nations. Sunstein also argues that legal resolution of constitutional questions, such as the once-common, now rejected, view that the U.S. Constitution permits racial segregation, may be the product of information cascades.

12.8 Group Adoption of Conventions

Social conventions often differ across groups and are sometimes transmitted from one group to the other. When there is limited observation across groups of agents, distinct cascades can temporarily form. Fisher and Wooders (2017) investigates the contagion of behavior between groups via an individual who belongs to both groups. In their model, two SBMs, labelled A and B, run in parallel. A single agent, I_n, $n > 2$, is common to both groups A and B and observes the actions of all predecessors in both groups. Every other agent is in only one group.

Suppose that n is odd, so that the difference between the number of Adopts and Rejects by I_n's predecessors is even. If neither group is in a cascade, then both groups cascade on I_n's action. If at least one group is in a cascade, then that cascade may be passed to the other group. If the two groups are in opposite cascades, then I_n follows her own signal, which breaks the cascade that opposes I_n's signal. If group A is in an Adopt cascade and group B is not in a cascade (i.e., the Adopt/Reject difference among I_n's predecessors in Group B is zero), then I_n Adopts and subsequently both groups are in an Adopt cascade, because I_n's successors in group B realize that I_n has superior social information. In summary, there can be contagion from the convention of one group to another group.

12.9 Belief Cascades

Several authors have suggested that the dynamics of public opinions, such as the rise and fall of ideologies and conspiracy theories, are information cascades.[58] One possible pathway might be the binary decision that people often face of whether or not to forward a news item on social media. Under appropriate conditions, information cascades in this decision are possible.

However, people on social media do not just forward items; they express opinions. This raises an important question which is not addressed in the basic cascades setting and the extensions we have considered so far: should we expect to see information cascades in publicly expressed *beliefs*?

In the SBM, agent I_n's belief $\hat{p}_n$ about the state is a continuous variable. Suppose that we extend the SBM by having each agent directly and truthfully communicate a report of her belief (probability assessment) to the next agent. Then each agent's private signal always influences the next agent's belief, so no agent is ever in a cascade of reported beliefs. However, in alternative settings there can be cascades—including incorrect cascades—in reported beliefs. We focus on one possible source of belief cascades, based upon discreteness in the communication channel.

[58] Sunstein (2019*a*, p.50) argues that the sudden popularity of "fake news" reports in recent years derives from information cascades. Kuran and Sunstein (1999, 2000) have also suggested that information cascades of belief play a role in how the public responds to health and financial crises, and resulting demands for risk regulation. Sunstein (2019*b*, p.18) describes the adoption of the belief that genetic modification of food is or is not a serious problem as an information cascade. He further argues that there are cascades in norms as well as in factual beliefs.

People often communicate in coarse categories, such as binary partitions. When asked about a model of car, people often say that they like or dislike it, as opposed to reporting intensities of liking on a continuum, detailed reasons for their opinions, or probability distributions over which model is the best.[59] Similarly, when asked for a suggestion about a movie or restaurant, or when discussing political topics, people often just name the preferred option, or say that an option is "hot" versus "sucks;" or "cool" versus "bogus." Such digitized communication is inevitable owing to limited time and cognitive resources of speakers and listeners.

To capture such limited bandwidth in the communication of beliefs, we consider a setting in which agents make decisions in a commonly known deterministic sequence, and each agent I_m observes the actions of some subset of the agent's predecessors.[60] Suppose that the action of each agent I_n is to report a binary indicator of state, $a_n \in \{L, H\}$.

We interpret the report a_n as the *reported belief* about the state. We will view reported belief L as a claim that the state is probably L, and reported belief H that the state is probably H.[61] This could take the form of expressing that "Candidate H will win the election" or "Candidate H will lose the election." Specifically, I_n *reports* belief L if her true belief $\hat{p}_n < 0.5$, reports H if $\hat{p}_n > 0.5$, and follows some indifference rule, such as flipping a coin between reporting L or H, if $\hat{p}_i = 0.5$. This reporting rule can be endogenized if agents desire a reputation with receivers for making good reports.

A *belief cascade* is defined as a situation in which, having received the reported

[59]As the well-known hedge fund manager and self-improvement author Ray Dalio put it, "It is common for conversations to consist of people sharing their conclusions rather than exploring the reasoning that led to those conclusions."

[60]Alternatively, each agent might also recall and pass on the history of all beliefs that have been reported to that agent. If signals are discrete, such a setting would effectively be equivalent to the model of BHW.

[61]A possible behavioral extension of the model would be to have the reports sometimes understood naively by receivers as indicating that the state is L or H for sure. If, with some probability, receivers make this mistake, information cascades can start very quickly.

beliefs of some set of predecessors, an agent's reported belief is independent of the agent's private signal. A belief cascade is also an information cascade, as defined earlier, since the reported belief a_n can be interpreted as the agent's action.

In this model, if signals are bounded, then as in § 4, asymptotic learning can fail; even in the limit, reported and actual beliefs can be incorrect. Furthermore, if signals are also discrete, then for reasons similar to those of the model of BHW, there are incorrect belief cascades (see footnote 60).

In the modified SBM example above, belief cascades derive from communication bandwidth constraints. However, belief cascades may also derive from other possible mechanisms. In models of message sending with payoff interactions, message senders sometimes strategically report only coarsened versions of their beliefs (Crawford and Sobel (1982)). This raises the possibility that sequential reporting might result in cascades, which is an interesting topic for future research.[62]

Whether belief cascades occur in practice is an empirical question. A possible application is to the spread of scientific claims. Since it is costly to read original source articles in detail, scholars often rely on descriptions of the original findings in later publications. Greenberg (2009) provides bibliometric evidence from the medical literature about the spread via chains of citations of unfounded scientific claims.

The broader message is that a rich direction for future research is the appli-

[62]Alternatively, in behavioral models with dual cognitive processing, an agent consists of two selves (sometimes called the "planner" and the "doer") who face distinct decision problems. This can correspond, for example, to the System 2 versus System 1 thinking distinction of Kahneman (2011). Suppose that the planner rationally updates and assesses probabilities, and, when $\hat{p}_i > 0.5$ instructs the doer that the state is H, and, when $\hat{p}_i < 0.5$ instructs the doer that the state is L. Such simplified instructions may be needed owing to limited cognitive processing power of the doer, who may face problems of distraction and time constraint. If people are typically in "doer" mode when engaged in casual conversation with others, then only the coarsened information is reported. This would again result in cascades in reported beliefs. Furthermore, these reported beliefs correspond to the genuine beliefs of people when in doer mode.

cation of social learning theory to the decision of which beliefs to publicly espouse. This may provide new insight into the spread of political and religious ideologies, folk economic thinking, and conspiracy theories.

13 Further Directions

We next discuss a number of possible avenues of future theoretical research on information cascades and social learning.

13.1 Information Design

A fruitful direction for further exploration is the problem of information design in sequential social learning settings. One possible purpose could be to improve general welfare by improving information aggregation. We have seen that this can sometimes be achieved by assigning some agents to be sacrificial lambs. Alternatively, a manager or a seller can design or influence the signal structure or the observation network to achieve private objectives. This amounts to a form of sequential Bayesian persuasion (Kamenica and Gentzkow (2011).

For example, the prices set by a monopolist influence the decisions of early buyers, which in turn affect the information and decisions of later buyers (Welch (1992)). Similarly, a principal can design an organization's observation structure to motivate agents to work hard or to choose desired actions (Bose et al. (2008), and Khanna and Slezak (2000) discussed in § 12.1).

In the papers in the preceding paragraph, the information that agents are endowed with is taken as given. There are also models wherein a principal derives information from observation of agents' payoffs and then conveys some of this information to other agents (Kremer, Mansour and Perry (2014); Che and Hörner (2018)). This contrasts with the papers discussed in § 12.1, in which a planner decides who observes whom.

As discussed in § 13.2, social learning modeling can potentially be extended to study social discourse, including the spread of publicly reported beliefs. The rise of the internet has given organizations much more flexibility to design social media platforms to influence who observes whom, and to curate what (mis)information an agents sees. This includes users' verbal opinions, recommendations, and anecdotes, as well as visual, aural and even physiological information. So the design of institutions to shape social learning is central to modern value creation and public policy. This has led to much debate about the design of social media algorithms.

Several questions about information design remain unanswered. How does the presence of bots and fake postings affect social learning? When is it beneficial for a monopolist to have superior information about the product? When do principals benefit from the ability to commit as to possible dependence of the observation structure on history? How does the design of private signal assignment or observation structure interact with other levers for influencing buyer beliefs, such as setting prices?[63] How do the interests of social media platform owners diverge from maximizing social welfare, and can regulatory policies address such divergence?

13.2 Imperfect Rationality

Turning to imperfect rationality in social learning, some models allow for overconfidence, or for agents misestimating the extent to which others are observing predecessors. In models with repeated actions on networks, the common assumption of myopia as a weakening of the assumption of forward-looking rationality (Gale and Kariv (2003*b*)).

[63]For example, a seller can send free samples of a product to influencers, effectively providing them with private signals about quality, or can pay a social media platform to inform users that certain of their friends have bought the product.

As discussed in § 6.2, biases that cause agents to be more aggressive in using their own signals tend to improve social information aggregation, whereas biases in the opposite direction worsen it. These findings suggest that it may be possible to derive conclusions about the relative success of individuals and groups that contain different distributions of psychological agent types (e.g., Bernardo and Welch (2001)).

The distribution of psychological types affects both whether there is asymptotic learning, and, when there is not, how much social learning occurs. In evolutionary settings, the success of individual types depends on a balance between the direct benefit of having a psychological trait versus the benefits to being in a group in which others have that trait. Analysis of this tension can endogenize the extent to which imperfectly rational types will survive. Such survival affects welfare, since, for example, types that heavily use their own private signals provide an informational benefit to other agents. The effectiveness of different kinds of policy interventions also depend on the biases of surviving agents.

Owing in part to different assumptions about the form of imperfect rationality, existing models (Bohren (2016), Bohren and Hauser (2019), Frick, Iijima and Ishii (2020)) obtain divergent conclusions about whether a small bias can have large effects. In some cases biases tend to be socially amplified, generating even larger and more persistent effects on social outcomes, such as failures of asymptotic learning. Whether small biases are amplified potentially has significant empirical and policy implications.

Another question is how imperfectly rational social learning influences boom/bust dynamics in the adoption of behaviors, such as investments, or mergers and acquisitions. We have discussed rational models of booms and busts (Chamley and Gale (1994)), including some within market settings (Lee (1998)). However, many observers have argued that imperfect rationality contributes to booms and busts, as in the models of Bohren and Hauser (2019) and Hirshleifer and Plotkin (2021). It will be valuable to develop both rational and imperfectly rational approaches in applications and empirical testing.

There is also extensive evidence that managers are imperfectly rational (see, e.g., Malmendier and Tate (2005)). A further promising direction is to explore product or security market settings with imperfectly rational social learning (see Eyster, Rabin and Vayanos (2018)).

An especially rich direction is to analyze how social learning and information cascades influence verbal discourse. This may provide insights into the growing concerns about the spread of misinformation and fake news. The wide availability of textual and social media databases (and associated empirical methods) in turn offers rich data to test hypotheses about social learning and misinformation. The model of belief cascades in § 12.9 (in which belief reports play the roles of actions) is a possible initial step. A further direction will be to incorporate a wider set of psychological biases that influence human communication. These may include neglect of selection in what others choose to assert, and confirmation bias. Such analysis can provide insight into why in some domains discourse grows ever more sophisticated (as in the scientific scholarly community), whereas in others discourse remains persistently naïve (as with much political discourse).

13.3 Repeated Actions

As discussed in § 11.3, the literature studying repeated actions by Bayesian agents is nascent. Only initial steps have been taken in the study of the speed of learning in such settings (see Harel et al. (2021)). Likewise, the question of which social networks promote more efficient information aggregation deserves further exploration (Golub and Jackson (2010); Mossel, Sly and Tamuz (2015))

When agents act repeatedly, new strategic incentives arise. An agent may seek to induce others to act in a way that will convey useful information. For example, consider an agent whose private information opposes the agent's social information. Even if following this private information lowers the agent's current payoff, doing so may cause others to take more informative actions. The topic of strategic informational incentives has remained almost completely unexplored (but see the “mad king” example in Mossel, Sly and Tamuz (2015)).

Only a handful of papers consider social learning with a changing state (see BHW, Moscarini, Ottaviani and Smith (1998); Frongillo, Schoenebeck and Tamuz (2011) and § 12.5 on fads); even fewer do so in a repeated action setting (Dasaratha, Golub and Hak (2019)). Many interesting questions remain unexplored. In a steady state, how well attuned are actions to state? How does this concordance depend on the private signal structure, or on deviations from rationality? How quickly does the preponderant action change after a change in state? Is the probability of change in the preponderant action less than or greater than the probability of a change in state (inertia or impulsiveness respectively; see Hirshleifer and Welch (2002) and Huang (2022))?

13.4 The Speed of Learning

The literature emphasizes the issue of whether or not asymptotic learning occurs, rather than how quickly learning occurs (a few recent exceptions are noted in § 4.3.5). Under a social welfare function with discounting, the welfare of earlier agents is especially important. So it is vital that learning occur quickly, not just asymptotically. Indeed, information externalities delay learning relative to the socially optimal rate. The question of how closely outcomes can approach an efficient benchmark deserves further exploration (see Chamley (2004*a*); Smith, Sørensen and Tian (2021)).

Social learning theory suggests a role for policy. Policymakers can potentially improve welfare by having agents with lower precisions act first, or by reducing the ability of early agents to observe information (whether from payoffs, public disclosures, or the actions of other agents). Such policies can encourage early agents to take informative actions, inducing greater information aggregation before the start of cascades (see also § 13.1).

A basic theme about social learning is the principle of countervailing adjustment (see § 6.1)—a shift in model parameters that promotes more accurate social information early on often tends to hinder the incorporation of new information socially. An issue for further exploration is the extent to which the principle of countervailing adjustment causes speed of learning to be somewhat insensitive to parameters of the learning process. Intuitively, parameter shifts that have a direct effect, at a given point in time, of making social information more accurate will tend to make the given agent's actions depend more weakly on her own private signal. This can oppose the direct increase in social information induced by the exogenous shift.

14 Conclusion and Empirical Testing

Social learning and information cascades have offered a new understanding of three phenomena. The first is that individuals who learn by observing others often converge over time to the same action. The second is that individuals who in the aggregate possess extensive information converge surprisingly often on an incorrect action, or only very slowly on the right one. Relatedly, there is path-dependence in social outcomes. The third is that people are attached to popular actions very loosely, resulting in fragility of mass outcomes, and fads.

In the social learning explanation, owing to an information externality, the information in the history of actions accumulates more slowly than would be socially optimal. In the case of information cascades, once a cascade starts, further information aggregation is completely blocked, resulting in social outcomes that are path-dependent and often incorrect. Furthermore, in fairly general settings, the amount of private information aggregated in the social belief is bounded above by the maximal information in just two private signals, as summarized in the Two-Signal Principle (§ 6.1). When there is poor information aggregation, even a small shock can cause actions to shift abruptly. The basic cascades model therefore implies that in such circumstances there will be occasional, sporadic bouts of sudden change. The system evolves towards a precarious position, instead of having fragility arise only under rare environmental conditions.

Information externalities in social learning, and incorrect cascades in particular, occur under reasonably mild assumptions. As we have discussed, these effects have been applied to a wide range of phenomena. The literature has provided a rich set of insights, e.g., by allowing for the possibilities that agents can choose when to act or whether to acquire private signals, or that the action being selected is a trade in a marketplace. Recent work has opened new vistas by analyzing social learning when agents take repeated actions, are imperfectly rational, or observe others in complex social networks. We highlight several directions for further research.

Our review has focused on theoretical models of social learning and information cascades. Although some of this wide-ranging literature is technical, the key insights are accessible for both applied theory and empirical research.

A general challenge for empirical testing of social learning theories is that conformity can derive from various sources. These include independent observations of the same information, a direct preference for conforming to the behavior of others, sanctions upon deviants, positive payoff externalities, and social learning—rational or otherwise.

In general, social influence (via social learning or other mechanisms, such as conformist preferences) cannot be derived solely from choice outcomes without information about causal relationships. If A and B behave similarly, this could be because A influences B, B influences A, or they are both subject to some other influence. Fortunately, there is often additional information to exploit about the order in which agents act (assuming an agent can only be influenced by someone who acted earlier) and the network of observation. Much of the literature on testing for social influence (often called "peer effects") examines conditions for identification in the absence of information about the sequencing of actions. Manski (1993) describes the "reflection problem," which is that in certain simple social influence settings, researchers cannot infer from the average behavior alone whether this average behavior has influenced the decisions of the individuals in a group. A further literature shows, however, that relatively modest restrictions on the structure of social influence allow identification of peer effects based on observations of individual behaviors, or even from average behaviors of groups (Blume et al. (2015)).

Turning to tests of sequential social learning, data with information about the order in which agents act provides valuable further restrictions. However, sequencing information alone does not necessarily fully identify social learning effects.[64]

Further possible partial identification derives from differences in the predictions of different mechanisms of the effects of aggregate shocks. For example, a key implication of the basic information cascades setting is that outcomes are fragile. This contrasts sharply with settings in which conformity derives from positive payoff externalities, positive preference interactions, or from sanctions upon deviants. In simple models based on conformity or externality effects, small shocks have small effects. In consequence, the observation of large shifts in aggregate behaviors without large shocks to the model parameters opposes such models (a possible example being the surgical fads and quack medical treatments mentioned in the introduction). In contrast, such observations are consistent with models based upon information cascades, which do imply fragility of aggregate outcomes.

More generally, three empirical approaches to distinguishing sources of homogeneous behavior are to:

(1) Make assumptions about causal pathways.

(2) Use auxiliary evidence about outcomes (such as the effects of external shocks or

[64]If I_2 and I_1 choose the same action, and we know that I_1 acted first (with no repeated moves), we can rule out that I_1 was influenced by I_2, but does rule out that they both may have been influenced by some common external information source. On the other hand, Blume et al. (2011) discuss a linear setting in which data on repeated choices in a population over time enables identification.

of repeated experiments) to rule out possible mechanisms.

(3) Obtain more granular data on the structure of private information signals, or on the observation network (such as the order in which agents act and observe others). Alternatively, design the observation network and signal structure in a laboratory or field experiment. In other words, exploit direct knowledge of exactly who saw what and when.

Approach (1) requires judgment about which potential theoretical explanations are plausible in the particular applied context. For example, the conclusion of Amihud, Hauser and Kirsh (2003) that there are information cascades in stock trading relies on the observation of the similarity in purchase decisions by small investors. The attribution of the outcome to social learning implicitly rules out other less plausible sources, such as traders gaining direct utility from imitating other traders, traders being pressured into imitating other traders by the threat of punishment, or traders having access to similar sources of information.

Also applying approach (1), Young (2009) considers three causes of diffusion of innovations: (i) local contagion wherein, much as in the spread of a disease, people adopt when they come in contact with others who have adopted; (ii) preference for conformity with aggregate choices in the population, rather than by contact with a local adopter and (iii) social learning. Unlike social learning, in Young's classification, contagion and influence are not information-based processes. Using additional model assumptions, he shows that the three mechanisms which lead to different types of adoption curves. Young finds that the S-shaped adoption curve in the seminal study of hybrid corn diffusion in Ryan and Gross (1943) is consistent with social learning, and not with local contagion nor preference for conformity.

As an example of approach (2), Salganik, Dodds and Watts (2006) provide evidence of conformity in song downloads in a field experiment on music marketplaces—some songs are downloaded much more than others (beyond what would be expected by chance). As opposed to social influence, a possible explanation is that some songs are just more appealing than others. However, the authors argue that this is unlikely to be the full explanation, because the correlations between song popularity rankings across marketplaces is low. Duan, Gu and Whinston (2009) find that downloads of free software at CNET decline precipitously when the download popularity rank decreases. They provide indirect identification in favor of information cascades by using further evidence to eliminate alternative competing hypotheses.

Often there are *negative* payoff externalities from following the same action, as with people joining a queue for the same restaurant or takeover bidders vying for the same target firm. In such contexts, observation of homogeneous behavior rules out payoff externalities as the dominant force, though it does not pinpoint social learning as the sole possible cause of conformity.

Alternatively, social learning can be distinguished from other proposed mechanisms by estimating directly whether the underpinnings of those proposed mechanisms are quantitatively important. For example, it may be possible to estimate the extent of payoff externalities between firms by examining how firms' profits are influenced by the actions of other firms (using appropriate instruments to identify causality).

An example of approach (3) is provided by Zhang (2010) using data on the order in which possible recipients of kidney for transplant are offered kidneys, and can conclude that predecessors did not accept the kidney. This knowledge about observation structure allows the inference that social learning has occurred.

Also under approach (3), experiments can be used to design signal and observation structures to test hypotheses and identify social learning effects, as in the experiments of Anderson and Holt (1997). They provide evidence that information cascades occur, but also provide evidence of deviations from rational behavior. Sometimes, inconsistent with rationality, subjects who should be in a cascade instead follow their own private signals.[65] In other experiments (Weizsäcker (2010), Ziegelmeyer,

[65]This is more common than the opposite mistake, wherein people who should not be in a cascade follow their immediate predecessors.

March and Krügel (2013)), agents with signals opposing predecessors systematically tend to fall into cascades only if the social information is non-trivially stronger than the agent's signal. A review of the experimental literature more generally indicates that subjects usually do not make sufficient use of social information (Morin et al. (2021)). This is consistent, for example, with overconfidence, as in Bernardo and Welch (2004). When agents over-rely on their private signals in this fashion, there tends to be more information aggregation before a cascade starts.

Several developments have created exciting opportunities for empirical testing for social learning effects. First, social media and other electronic databases have greatly increased the information available to empiricists about the network of social observation and connectedness in relation to behavior. New network datasets include the publicly available Facebook Social Connectedness Index[66], and are being applied to measure social influence—and in some cases social learning effects—on economic decisions (Bailey et al. (2018*b*), Bailey et al. (2022), Bali et al. (2022)).

Second, textual and in some cases social network data provide possible proxies for the information signals available to different agents or the beliefs that such agents have formed. This approach has been used in a literature on the effects of linguistic tone on economic outcomes (Huang, Teoh and Zhang (2014)). Potentially complementing the use of textual data to measure the private information signals of agents or the information that is conveyed to others, there are also new methods of extracting quantitative information from other aspects of communication, such as vocal tone (Mayew and Venkatachalam (2012)) and facial expression in conversation (Flam, Green and Sharp (2020)), and the use of images (Nekrasov, Teoh and Wu (2021)) and video (Gu, Teoh and Wu (2022)) in electronic communications. Third, the application of laboratory experimentation to social learning has provided more sharply controlled tests, and field experimentation promises to address concerns about external validity.

[66] See, for example, `https://github.com/social-connectedness-index`, `https://data.humdata.org/dataset/social-connectedness-index`, and Bailey et al. (2018*a*).

For example, there is increasing use of field experiments to study social learning in various applications (Banerjee et al. (2013), Cai, Janvry and Sadoulet (2015), Banerjee et al. (2019*a*), Breza and Chandrasekhar (2019), and Banerjee et al. (2019*b*)).

For all these reasons, empirical testing is playing an increasingly important role in the further advance of this field—both in testing existing models and in stimulating new theorizing. An interesting direction for further exploration is whether conformists or deviants from the predominant action make better decisions. In some empirical settings, agents' actions and payoffs are indeed observable to the researcher. The answer to this question differs across models, so addressing it empirically, in conjunction with refinement of existing models, may help distinguish alternative mechanisms of social influence.

Indeed, we close this review by observing that two key advances provide opportunities for synergistic advance. First is documenting empirically and incorporating in our models how agents are embedded in social observation networks. Second is documenting and modeling the ways in which agents are imperfectly rational. One example is the tendency of agents to rely on their private signals even when rationally it pays to imitate, as discussed above.

Both advances are being supported by the new experimental testing approaches and the rise of rich new network and textual datasets described above. We expect this interplay between these advances and theoretical modeling to be increasingly productive in generating new insights into social learning and information cascades.

References

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar. 2019. "Learning From Reviews: The Selection Effect and the Speed of Learning." Working Paper, Massachusetts Institute of Technology.

Acemoglu, Daron, and Asuman E. Ozdaglar. 2011. "Opinion Dynamics and Learning in Social Networks." *Dynamic Games and Applications*, 1(1): 3–49.

Acemoglu, Daron, and James A. Robinson. 2006. *Economic Origins of Dictatorship and Democracy.* New York, NY:Cambridge University Press.

Acemoglu, Daron, Munther Dahleh, Ilan Lobel, and Asuman E. Ozdaglar. 2011. "Bayesian Learning in Social Networks." *Review of Economic Studies*, 78(4): 1201–1236.

Aghamolla, Cyrus. 2018. "Observational Learning, Endogenous Timing, and Information Acquisition." Working Paper, University of Minnesota.

Aghamolla, Cyrus, and Ilan Guttman. 2021. "Strategic timing of IPOs and disclosure: A dynamic model of multiple firms." *Accounting Review*, 96(3): 27–57.

Aghamolla, Cyrus, and Tadashi Hashimoto. 2020. "Information arrival, delay, and clustering in financial markets with dynamic freeriding." *Journal of Financial Economics*, 138(1): 27–52.

Akerlof, George A. 1976. "The Economics of Caste and of the Rat Race and Other Woeful Tales." *Quarterly Journal of Economics*, 90(4): 599–617.

Alevy, Jonathan E., Michael S. Haigh, and John A. List. 2007. "Information Cascades: Evidence from a Field Experiment with Financial Market Professionals." *Journal of Finance*, 62(1): 151–180.

Ali, S. Nageeb. 2018*a*. "Herding with costly information." *Journal of Economic Theory*, 175: 713–729.

Ali, S. Nageeb. 2018*b*. "On the Role of Responsiveness in Rational Herds." *Economics Letters*, 163: 79–82.

Ali, S. Nageeb, and Navin Kartik. 2012. "Herding with collective preferences." *Economic Theory*, 51(3): 601–626.

Ali, S Nageeb, Jacob K Goeree, Navin Kartik, and Thomas R Palfrey. 2008. "Information aggregation in standing and ad hoc committees." *American Economic Review*, 98(2): 181–86.

Amihud, Yakov, Shmuel Hauser, and Amir Kirsh. 2003. "Allocations, adverse selection, and cascades in IPOs: Evidence from the Tel Aviv Stock Exchange." *Journal of Financial Economics*, 68(1): 137–158.

Anderson, Lisa R., and Charles A. Holt. 1997. "Information Cascades in the Laboratory." *American Economic Review*, 87(5): 847–862.

Anderson, Lisa R., and Charles A. Holt. 2008. "Information Cascade Experiments." In *Handbook of Experimental Economics Results*. Vol. 1, , ed. Charles R. Plott and Vernon L. Smith, Chapter 39, 335–343. Elsevier.

Anderson, Ryan C., and Michele K. Surbey. 2014. "I Want What She's Having: Evidence of Human Mate Copying." *Human Nature*, 25(3): 342–358.

Arieli, Itai, and Manuel Mueller-Frank. 2019. "Multidimensional Social Learning." *Review of Economic Studies*, 86(3): 913–940.

Arthur, W. Brian. 1989. "Competing Technologies, Increasing Returns, and Lock-in by Historical Events." *Economic Journal*, 99(394): 116–131.

Arya, Anil, Jonathan Glover, and Brian Mittendorf. 2006. "Hierarchical Reporting, Aggregation, and Information Cascades." *Managerial and Decision Economics*, 27(5): 355–362.

Aumann, Robert J. 1976. "Agreeing to Disagree." *Annals of Statistics*, 4(6): 1236–1239.

Avery, Christopher, and Peter Zemsky. 1998. "Multi-Dimensional Uncertainty and Herd Behavior in Financial Markets." *American Economic Review*, 88(4): 724–48.

Bailey, Michael, Drew M. Johnston, Theresa Kuchler, Johannes Stroebel, and Arlene Wong. 2022. "Peer Effects in Product Adoption." *American Economic Journal: Applied Economics*. Forthcoming.

Bailey, Michael, Rachel Cao, Theresa Kuchler, Johannes Stroebel, and Arlene Wong. 2018*a*. "Social Connectedness: Measurement, Determinants, and Effects." *Journal of Economic Perspectives*, 32(3): 259–280.

Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel. 2018*b*. "The Economic Effects of Social Networks: Evidence from the Housing Market." *Journal of Political Economy*, 126(6): 2224–2276.

Bala, Venkatesh, and Sanjeev Goyal. 1998. "Learning from Neighbours." *Review of Economic Studies*, 65(3): 595–621.

Bali, Turan G., David A. Hirshleifer, Lin Peng, and Yi Tang. 2022. "Attention, Social Interaction, and Investor Attraction to Lottery Stocks." Working paper, Georgetown University.

Banerjee, Abhijit. 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, 107(3): 797–817.

Banerjee, Abhijit, and Drew Fudenberg. 2004. "Word of Mouth learning." *Games and Economic Behavior*, 46(1): 1–22.

Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. 2019*a*. "Using gossips to spread information: Theory and evidence from two randomized controlled trials." *Review of Economic Studies*, 86(6): 2453–2490.

Banerjee, Abhijit, Emily Breza, Arun G. Chandrasekhar, and Benjamin Golub. 2019*b*. "When less is more: Experimental evidence on information delivery during India's demonetization." Working Paper, Massachusetts Institute of Technology.

Banerjee, Abhijit V., Arun Chandrasekhar, Esther Duflo, and Matthew O. Jackson. 2013. "The Diffusion of Microfinance." *Science*, 341(6144): 1236498.

Bartels, Larry M. 1988. *Presidential Primaries and the Dynamics of Public Choice.* Princeton, New Jersey:Princeton University Press.

Battaglini, Marco. 2005. "Sequential voting with abstention." *Games and Economic Behavior*, 51(2): 445–463.

Bernardo, Antonio, and Ivo Welch. 2001. "On the Evolution of Overconfidence and Entrepreneurs." *Journal of Economics and Management Strategy*, 10(3): 301–330.

Bernardo, Antonio E., and Ivo Welch. 2004. "Liquidity and Financial Market Runs." *Quarterly Journal of Economics*, 119(1): 135–158.

Bernheim, B. Douglas. 1994. "A Theory of Conformity." *Journal of Political Economy*, 102(5): 841–877.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy*, 100(5): 992–1026.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1998. "Learning from the behavior of others: Conformity, fads, and informational cascades." *Journal of Economic Perspectives*, 12(3): 151–170.

Blume, Lawrence E. 1993. "The statistical mechanics of strategic interaction." *Games and economic behavior*, 5(3): 387–424.

Blume, Lawrence E., William A. Brock, Steven N. Durlauf, and Rajshri Jayaraman. 2015. "Linear social interactions models." *Journal of Political Economy*, 123(2): 444–496.

Blume, Lawrence E., William A. Brock, Steven N. Durlauf, and Yannis M. Ioannides. 2011. "Identification of Social Interactions." In *Handbook of Social Economics*. Vol. 1, Chapter 18, 853–964. Amsterdam: North Holland.

Bobkova, Nina, and Helene Mass. 2020. "Learning What Unites or Divides Us: Information Acquisition in Social Learning." Working paper, Rice University.

Boccaletti, Stefano, Vito Latora, Yamir Moreno, Martin Chavez, and D-U Hwang. 2006. "Complex networks: Structure and Dynamics." *Physics Reports*, 424(4-5): 175–308.

Bohren, J. Aislinn. 2016. "Informational herding with model misspecification." *Journal of Economic Theory*, 163: 222–247.

Bohren, J. Aislinn, and Daniel N. Hauser. 2019. "Social Learning with Model Misspecification: A Framework and a Robustness Result." Working paper, University of Pennsylvania.

Bolton, Patrick, and Christopher Harris. 1999. "Strategic Experimentation." *Econometrica*, 67(2): 349–374.

Bonacich, Phillip. 1987. "Power and centrality: A family of measures." *American Journal of Sociology*, 92(5): 1170–1182.

Bose, Subir, Gerhard Orosel, Marco Ottaviani, and Lise Vesterlund. 2008. "Monopoly pricing in the binary herding model." *Economic Theory*, 37(2): 203–241.

Bowers, Robert I., Skyler S. Place, Peter M. Todd, Lars Penke, and Jens B. Asendorpf. 2012. "Generalization in mate-choice copying in humans." *Behavioral Ecology*, 23(1): 112–124.

Boyd, Robert, and Peter J. Richerson. 1985. *Culture and the Evolutionary Process.* Chicago, IL:University of Chicago Press.

Breza, Emily, and Arun G. Chandrasekhar. 2019. "Social networks, reputation, and commitment: Evidence from a savings monitors experiment." *Econometrica*, 87(1): 175–216.

Brunnermeier, Markus. 2001. *Asset Pricing under Asymmetric Information: Bubbles, Crashes, Technical Analysis and Herding.* Oxford, UK:Oxford University Press.

Burguet, Roberto, and Xavier Vives. 2000. "Social Learning and Costly Information Acquisition." *Journal of Economic Theory*, 15(1): 185–205.

Cai, Jing, Alain De Janvry, and Elisabeth Sadoulet. 2015. "Social networks and the decision to insure." *American Economic Journal: Applied Economics*, 7(2): 81–108.

Callander, Steven. 2007. "Bandwagons and Momentum in Sequential Voting." *Review of Economic Studies*, 74: 653–684.

Callander, Steven, and Johannes Hörner. 2009. "The wisdom of the minority." *Journal of Economic Theory*, 144(4): 1421–1439.

Caminal, Ramon, and Xavier Vives. 1996. "Why Market Shares Matter: An Information-Based Theory." *The RAND Journal of Economics*, 27(2): 221–239.

Caminal, Ramon, and Xavier Vives. 1999. "Price Dynamics and Consumer Learning." *Journal of Economics and Management Strategy*, 8(1): 95–131.

Cao, H. Henry, and David Hirshleifer. 1997. "Limited observability, reporting biases, and informational cascades." Working paper, Fisher College of Business, The Ohio State University.

Cao, H. Henry, Bing Han, and David Hirshleifer. 2011. "Taking the Road Less Traveled by: Does Conversation Eradicate Pernicious Cascades?" *Journal of Economic Theory*, 146(4): 1418–1436.

Caplin, Andrew, and John Leahy. 1993. "Sectoral Shocks, Learning and Aggregate Fluctuations." *Review of Economic Studies*, 60: 777–794.

Caplin, Andrew, and John Leahy. 1994. "Business as Usual, Market Crashes, and Wisdom after the fact." *American Economic Review*, 84(3): 548–565.

Castellano, Claudio, Santo Fortunato, and Vittorio Loreto. 2009. "Statistical physics of social dynamics." *Reviews of Modern Physics*, 81(2): 591.

Celen, Bogachan, and Shachar Kariv. 2004*a*. "Distinguishing Informational Cascades from Herd Behavior in the Laboratory." *American Economic Review*, 94(3): 484–497.

Çelen, Boğaçhan, and Shachar Kariv. 2004*b*. "Observational learning under imperfect information." *Games and Economic Behavior*, 47(1): 72–86.

Chamley, Christophe, and Douglas Gale. 1994. "Information Revelation and Strategic Delay in Irreversible Decisions." *Econometrica*, 62(5): 1065–1085.

Chamley, Christophe P. 2004*a*. "Delays and equilibria with large and small information in social learning." *European Economic Review*, 48(3): 477–501.

Chamley, Christophe P. 2004*b*. *Rational Herds: Economic Models of Social Learning.* Cambridge, UK:Cambridge University Press.

Chari, V.V., and Patrick J. Kehoe. 2004. "Financial Crises as Herds: Overturning the critiques." *Journal of Economic Theory*, 119(1): 128–150.

Cheng, Yu, Wade Hann-Caruthers, and Omer Tamuz. 2018. "A deterministic protocol for sequential asymptotic learning." 1735–1738, IEEE.

Chen, Jaden Yang. 2021. "Sequential Learning under Informational Ambiguity." *Available at SSRN 3480231*.

Che, Yeon-Koo, and Johannes Hörner. 2018. "Recommender systems as mechanisms for social learning." *Quarterly Journal of Economics*, 133(2): 871–925.

Cipriani, Marco, and Antonio Guarino. 2008*a*. "Herd Behavior and Contagion in Financial Markets." *B.E. Journal of Theoretical Economics*, 8(1): Article 24.

Cipriani, Marco, and Antonio Guarino. 2008*b*. "Transaction Costs and Informational Cascades in Financial Markets." *Journal of Economic Behavior and Organization*, 68: 581–592.

Cipriani, Marco, and Antonio Guarino. 2014. "Estimating a Structural Model of Herd Behavior in Financial Markets." *American Economic Review*, 104(1): 224–251.

Crawford, Vincent P., and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica*, 50(6): 1431–1451.

Dall, Sasha R. X., Luc-Alain Giraldeau, Ola Olsson, John M. McNamara, and David W. Stephens. 2005. "Information and its use by animals in evolutionary ecology." *Trends in Ecology & Evolution*, 20(4): 187–193.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam. 1998. "Investor Psychology and Security Market Under- and Overreactions." *Journal of Finance*, 53(6): 1839–1885.

Dasaratha, Krishna, and Kevin He. 2019. "Aggregative Efficiency of Bayesian Learning in Networks." *arXiv preprint arXiv:1911.10116*.

Dasaratha, Krishna, Benjamin Golub, and Nir Hak. 2019. "Social Learning in a Dynamic Environment." Working paper, Harvard University.

Daughety, Andrew F., and Jennifer F. Reinganum. 1999. "Stampede to Judgment: Persuasive Influence and Herding Behavior by Courts." *American Law and Economics Review*, 1(1 & 2).

Décamps, Jean-Paul, and Stefano Lovo. 2006. "Informational cascades with endogenous prices: The role of risk aversion." *Journal of Mathematical Economics*, 42(1): 109–120.

Deffuant, Guillaume, David Neau, Frederic Amblard, and Gérard Weisbuch. 2000. "Mixing beliefs among interacting agents." *Advances in Complex Systems*, 3(01n04): 87–98.

DeGroot, Morris H. 1974. "Reaching a Consensus." *Journal of the American Statistical Association*, 69(345): 118–121.

Dekel, Eddie, and Michele Piccioni. 2000. "Sequential Voting Procedures in Symmetric Binary Elections." *Journal of Political Economy*, 108(1): 34–44.

Dekel, Eddie, Drew Fudenberg, and David K Levine. 2004. "Learning to play Bayesian games." *Games and Economic Behavior*, 46(2): 282–303.

DeMarzo, Peter, Dimitri Vayanos, and Jeff Zwiebel. 2003. "Persuasion Bias, Social Influence, and Uni-Dimensional Opinions." *Quarterly Journal of Economics*, 118: 909–968.

Duan, Wenjing, Bin Gu, and Andrew B. Whinston. 2009. "Informational Cascades and Software Adoption on the Internet: An Empirical Investigation." *MIS Quarterly*, 33(1): 23–48.

Dugatkin, Lee A., and Jean-Guy J. Godin. 1992. "Reversal of female mate choice by copying in the guppy (*Poecilia reticulata*)." *Proceedings of the Royal Society of London, Series B*, 249: 179–184.

Durlauf, Steven N, and Yannis M Ioannides. 2010. "Social interactions." *Annu. Rev. Econ.*, 2(1): 451–478.

Ellis, Christopher J., and John Fender. 2011. "Information Cascades and Revolutionary Regime Transitions." *Economic Journal*, 121(1): 763–792.

Ellison, Glenn. 1993. "Learning, local interaction, and coordination." *Econometrica: Journal of the Econometric Society*, 1047–1071.

Ellison, Glenn, and Drew Fudenberg. 1993. "Rules of thumb for social learning." *Journal of Political Economy*, 101(4): 612–643.

Ellison, Glenn, and Drew Fudenberg. 1995. "Word-of-Mouth Communication and Social Learning." *Quarterly Journal of Economics*, 110(1): 93–125.

Enke, Benjamin, and Florian Zimmermann. 2019. "Correlation Neglect in Belief Formation." *Review of Economic Studies*, 86(1): 313–332.

Epstein, Gil S. 2010. "Informational Cascades and the Decision to Migrate." In *Frontiers of Economics and Globalization*. , ed. Hamid Beladi and E. Kwan Choi, Chapter 2. Bingley, UK:Emerald Group Publishing Limited.

Eyster, Erik, and Matthew Rabin. 2010. "Naïve Herding in Rich-Information Settings." *American Economic Journal: Microeconomics*, 2(4): 221–243.

Eyster, Erik, and Matthew Rabin. 2014. "Extensive Imitation is Irrational and Harmful." *Quarterly Journal of Economics*, 129(4): 1861–1898.

Eyster, Erik, Andrea Galeotti, Navin Kartik, and Matthew Rabin. 2014. "Congested observational learning." *Games and Economic Behavior*, 87: 519–538.

Eyster, Erik, Matthew Rabin, and Dimitri Vayanos. 2018. "Financial markets where traders neglect the informational content of prices." *Journal of Finance*, 74(1): 371–399.

Fey, Mark. 2000. "Informational Cascades and Sequential Voting." *University of Rochester. Mimeo*.

Fisher, James C.D., and John Wooders. 2017. "Interacting Information Cascades: On the Movement of Conventions Between Groups." *Economic Theory*, 63(1): 211–231.

Flam, Rachel W., Jeremiah Green, and Nathan Y. Sharp. 2020. "Do Investors Respond to CEO Facial Expressions of Anger During Television Interviews?" *https://ssrn.com/abstract=3740755*.

Frick, Mira, Ryota Iijima, and Yuhta Ishii. 2020. "Misinterpreting Others and the Fragility of Social Learning." *Econometrica*, 88(6): 2281–2328.

Frongillo, Rafael M, Grant Schoenebeck, and Omer Tamuz. 2011. "Social learning in a changing world." 146–157, Springer.

Gale, Douglas. 1996. "What have we learned from social learning?" *European Economic Review*, 40: 617–628.

Gale, Douglas, and Shachar Kariv. 2003*a*. "Bayesian learning in social networks." *Games and Economic Behavior*, 45(2): 329–346.

Gale, Douglas, and Shachar Kariv. 2003*b*. "Bayesian Learning in Social Networks." *Games and Economic Behavior*, 45(2): 329–346.

Galeotti, Andrea, and Sanjeev Goyal. 2010. "The Law of the Few." *American Economic Review*, 100(4): 1468–1492.

Geanakoplos, John D., and Herakles M. Polemarchakis. 1982. "We can't disagree forever." *Journal of Economic Theory*, 192–200.

Gibson, Robert M., and Jacob Hoglund. 1992. "Copying and Sexual Selection." *Trends in Evolution and Ecology*, 7(7): 229–232.

Gibson, Robert M., Jack W. Bradbury, and Sandra L. Vehrencamp. 1991. "Mate Choice in Lekking Sage Grouse Revisited: The Roles of Vocal Display, Female Site Fidelity, and Copying." *Behavioral Ecology*, 2(2): 165–180.

Giraldeau, Luc-Alain, Thomas J. Valone, and Jennifer J. Templeton. 2002. "Potential disadvantages of using socially acquired information." *Philosophical Transactions: Biological Sciences*, 357(1427): 1559–1566.

Glosten, Lawrence R., and Paul R. Milgrom. 1985. "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders." *Journal of Financial Economics*, 14(1): 71–100.

Goeree, Jacob, Thomas Palfrey, and Brian Rogers. 2006. "Social learning with private and common values." *Economic Theory*, 28(2): 245–264.

Golub, Benjamin, and Evan Sadler. 2016. "Learning in Social Networks." In *The Oxford Handbook of the Economics of Networks*. , ed. Yann Bramoullé, Andrea Galeotti and Brian W. Rogers, Chapter 19, 504–542. New York, NY:Oxford University Press.

Golub, Benjamin, and Matthew O Jackson. 2010. "Naive learning in social networks and the wisdom of crowds." *American Economic Journal: Microeconomics*, 2(1): 112–49.

Greenberg, Steven A. 2009. "How citation distortions create unfounded authority: Analysis of a citation network." *British Medical Journal*, 339: b2680.

Grenadier, Steven R. 1999. "Information revelation through option exercise." *Review of Financial Studies*, 12(1): 95–129.

Grossman, Sanford J., and Joseph E. Stiglitz. 1976. "Information and Competitive Price Systems." *American Economic Review*, 66(2): 246–253.

Guarino, Antonio, and Philippe Jehiel. 2013. "Social Learning with Coarse Inference." *American Economic Journal: Microeconomics*, 5(1): 147–174.

Guarino, Antonio, Heike Harmgart, and Steffen Huck. 2011. "Aggregate Information Cascades." *Games and Economic Behavior*, 73(1): 167–185.

Gul, Faruk, and Russell Lundholm. 1995. "Endogenous Timing and Clustering of Agents' Decisions." *Journal of Political Economy*, 103(5): 1039–1066.

Gu, Ming, Siew Hong Teoh, and Shijia Wu. 2022. "Contagion of Investor Sentiment in Online Investment Communities: Evidence from Dynamic Visuals on Stocktwits." Working Paper, University of California, Irvine.

Hann-Caruthers, Wade, Vadim V. Martynov, and Omer Tamuz. 2018. "The speed of sequential asymptotic learning." *Journal of Economic Theory*, 173: 383–409.

Harel, Matan, Elchanan Mossel, Philipp Strack, and Omer Tamuz. 2021. "Rational Groupthink." *Quarterly Journal of Economics*, 136(1): 621–668.

Harris, Lawrence. 1991. "Stock price clustering and discreteness." *Review of Financial Studies*, 4(3): 389–415.

Hendricks, Kenneth, Alan Sorensen, and Thomas Wiseman. 2012. "Observational Learning and Demand for Search Goods." *American Economic Journal: Microeconomics*, 4(1): 1–31.

Hendricks, Kenneth, and Dan Kovenock. 1989. "Asymmetric Information, Information Externalities, and Efficiency: The Case of Oil Exploration." *Rand Journal of Economics*, 20(2): 164–182.

Henrich, Joseph, and Francisco Gil-White. 2001. "The Evolution of Prestige: Freely conferred status as a mechanism for enhancing the benefits of cultural transmission." *Evolution and Human Behavior*, 22(3): 165–196.

Herrera, Helios, and Johannes Horner. 2013. "Biased Social Learning." *Games and Economic Behavior*, 80: 131–146.

Hirshleifer, David, and Ivo Welch. 2002. "An Economic Approach to the Psychology of Change: Amnesia, Inertia, and Impulsiveness." *Journal of Economics and Management Strategy*, 11: 379–421.

Hirshleifer, David, and Joshua Plotkin. 2021. "Moonshots, Investment Booms, and Selection Bias in the Transmission of Cultural Traits." *Proceedings of the National Academy of Sciences*, 118(e2015571118).

Hirshleifer, David, and Siew Hong Teoh. 2003. "Herd Behavior and Cascading in Capital Markets: A Review and Synthesis." *European Financial Management*, 9(1): 25–66.

Hirshleifer, David, and Siew Hong Teoh. 2009. "Thought and Behavior Contagion in Capital Markets." In *Handbook of Financial Markets: Dynamics and Evolution*. *Handbooks in Finance*, , ed. Thorsten Hens and Klaus Schenk-Hoppe, Chapter 1, 1–46. Amsterdam, The Netherlands:North-Holland.

Hirshleifer, David, Yaron Levi, Ben Lourie, and Siew Hong Teoh. 2019. "Decision Fatigue and Heuristic Analyst Forecasts." *Journal of Financial Economics*, 133(1): 83–98.

Holley, Richard A, and Thomas M Liggett. 1975. "Ergodic theorems for weakly interacting infinite systems and the voter model." *Annals of Probability*, 643–663.

Horst, Ulrich, and Jose A Scheinkman. 2006. "Equilibria in systems of social interactions." *Journal of Economic Theory*, 130(1): 44–77.

Huang, Wanying. 2022. "The Emergence of Fads in a Changing World." Working Paper, California Institute of Technology.

Huang, Xuan, Siew Hong Teoh, and Yinglei Zhang. 2014. "Tone management." *Accounting Review*, 89(3): 1083–1113.

Ifrach, Bar, Costis Maglaras, Marco Scarsini, and Anna Zseleva. 2019. "Bayesian social learning from consumer reviews." *Operations Research*, 67(5): 1209–1221.

Ikenberry, David L., and James Peter Weston. 2008. "Clustering in US Stock Prices after Decimalisation." *European Financial Management*, 14(1): 30–54.

Jackson, Matthew O., and Ehud Kalai. 1997. "Social learning in recurring games." *Games and Economic Behavior*, 21(1-2): 102–134.

Janis, Irving L. 1982. *Groupthink: Psychological Studies of Policy Decisions and Fiascoes.* . 2nd ed., Boston:Houghton Mifflin.

Janis, Irving L., and Leon Mann. 1977. *Decision Making: A Psychological Analysis of Conflict, Choice, and Commitment.* New York:Free Press.

Kahn, Charles, George Pennacchi, and Ben Sopranzetti. 1999. "Bank Deposit Rate Clustering: Theory and Empirical Evidence." *Journal of Finance*, 54(6): 2185–2214.

Kahneman, Daniel. 2011. *Thinking, Fast and Slow.* New York, NY:Farrar, Straus and Giroux.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review*, 101(6): 2590–2615.

Kandori, Michihiro, George J Mailath, and Rafael Rob. 1993. "Learning, mutation, and long run equilibria in games." *Econometrica: Journal of the Econometric Society*, 29–56.

Kartik, Navin, SangMok Lee, Tianhao Liu, and Daniel Rappoport. 2022. "Observational Learning with Ordered States." Columbia.

Katz, Leo. 1953. "A new status index derived from sociometric analysis." *Psychometrika*, 18(1): 39–43.

Khanna, Naveen, and Steve L. Slezak. 2000. "The Effect of Organizational Form on Information Aggregation and Project Choice: The Problem of Informational Cascades." *Journal of Economics and Management Strategy*, 9(1): 115–156.

Kirman, Alan. 1993. "Ants, Rationality, and Recruitment." *Quarterly Journal of Economics*, 108(1): 137–156.

Knight, Brian, and Nathan Schiff. 2010. "Momentum and Social Learning in Presidential Primaries." *Journal of Political Economy*, 118(6): 1110–1150.

Kremer, Ilan, Yishay Mansour, and Motty Perry. 2014. "Implementing the "wisdom of the crowd"." *Journal of Political Economy*, 122(5): 988–1012.

Kreps, David M, Paul Milgrom, John Roberts, and Robert Wilson. 1982. "Rational cooperation in the finitely repeated prisoners' dilemma." *Journal of Economic theory*, 27(2): 245–252.

Kroft, Kory, Fabian Lange, and Matthew J. Notowidigdo. 2013. “Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment.” *Quarterly Journal of Economics*, 128(3): 1123–1167.

Kübler, Dorothea, and Georg Weizsäcker. 2003. “Information Cascades on the Labor Market.” *Journal of Economics*, 80(3): 211–229.

Kultti, Klaus K., and Paavo A. Miettinen. 2007. “Herding with costly observation.” *The BE Journal of Theoretical Economics*, 7(1): 1–16.

Kuran, Timur. 1987. “Preference Falsification, Policy Continuity and Collective Conservatism.” *Economic Journal*, 97(387): 642–665.

Kuran, Timur. 1989. “Sparks and Prairie Fires: A Theory of Unanticipated Political Revolution.” *Public Choice*, 61(1): 41–74.

Kuran, Timur, and Cass R. Sunstein. 2000. “Controlling availability cascades.” In *Behavioral Law and Economics*. , ed. Cass R. Sunstein, 374–397. Cambridge, UK:Cambridge University Press.

Kuran, Timur, and Cass Sunstein. 1999. “Availability Cascades and Risk Regulation.” *Stanford Law Review*, 51: 683–768.

Lee, In Ho. 1993. “On the convergence of informational cascades.” *Journal of Economic Theory*, 61: 395–411.

Lee, In Ho. 1998. “Market Crashes and Informational Avalanches.” *Review of Economic Studies*, 65: 741–759.

Le, Tho Ngoc, Vijay Subramanian, and Randall Berry. 2016. “Quantifying the Utility of Noisy Reviews in Stopping Information Cascades.” Submitted to the 55rd IEEE Conference on Decision and Control.

Liebowitz, Stan J, and Stephen E Margolis. 1995. “Path dependence, lock-in, and history.” *Journal of Law, Economics, & Organization*, 205–226.

Little, Anthony C., Robert P. Burriss, Benedict C. Jones, Lisa M.DeBruine, and Christine A. Caldwell. 2008. “Social influence in human face preference: men and women are influenced more for long-term than short-term attractiveness decisions.” *Evolution and Human Behavior*, 29(2): 140–146.

Lobel, Ilan, and Evan Sadler. 2015. “Information diffusion in networks through social learning.” *Theoretical Economics*, 10(3): 807–851.

Lohmann, Susanne. 1994. "The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, East Germany, 1989-91." *World Politics*, 47: 42–101.

Malmendier, Ulrike, and Geoffrey Tate. 2005. "Does Overconfidence Affect Corporate Investment? CEO Overconfidence Measures Revisited." *European Financial Management*, 11(5).

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies*, 60(3): 531–542.

Mayew, William J., and Mohan Venkatachalam. 2012. "The Power of Voice: Managerial Affective States and Future Firm Performance." *Journal of Finance*, 67(1): 1–43.

Milgram, Stanley, Leonard Bickman, and Lawrence Berkowitz. 1969. "Note on the drawing power of crowds of different size." *Journal of Personality and Social Psychology*, 13(2): 79–82.

Molavi, Pooya, Alireza Tahbaz-Salehi, and Ali Jadbabaie. 2018. "A Theory of Non-Bayesian Social Learning." *Econometrica*, 86(2): 445–490.

Monzón, Ignacio, and Michael Rapp. 2014. "Observational learning with position uncertainty." *Journal of Economic Theory*, 154: 375–402.

Morin, Olivier, Pierre Olivier Jacquet, Krist Vaesen, and Alberto Acerbi. 2021. "Social information use and social information waste." *Philosophical Transactions of the Royal Society B*, 376(1828): 20200052.

Moscarini, Giuseppe, Marco Ottaviani, and Lones Smith. 1998. "Social Learning in a Changing World." *Economic Theory*, 11(3): 657–665.

Mossel, Elchanan, Allan Sly, and Omer Tamuz. 2014. "Asymptotic learning on bayesian social networks." *Probability Theory and Related Fields*, 158(1-2): 127–157.

Mossel, Elchanan, Allan Sly, and Omer Tamuz. 2015. "Strategic Learning and the Topology of Social Networks." *Econometrica*, 83(5): 1755–1794.

Mossel, Elchanan, Manuel Mueller-Frank, Allan Sly, and Omer Tamuz. 2020. "Social learning equilibria." *Econometrica*, 88(3): 1235–1267.

Mueller-Frank, Manuel. 2014. "Does one Bayesian make a Difference?" *Journal of Economic Theory*, 154: 423–452.

Mueller-Frank, Manuel, and Mallesh M. Pai. 2016. "Social Learning with Costly Search." *American Economic Journal: Microeconomics*, 8(1): 83–109.

Nekrasov, Alexander, Siew Hong Teoh, and Shijia Wu. 2021. "Visuals and Attention to Earnings News on Twitter." *Review of Accounting Studies*, 1–43. Forthcoming.

Newberry, Peter W. 2016. "An empirical study of observational learning." *RAND Journal of Economics*, 47(2): 394–432.

Oberholzer-Gee, Felix. 2008. "Nonemployment stigma as rational herding: A field experiment." *Journal of Economic Behavior & Organization*, 65(1): 30–40.

Ottaviani, Marco, and Peter Norman Sørensen. 2000. "Herd behavior and investment: Comment." *American Economic Review*, 90(3): 695–704.

Ottaviani, Marco, and Peter Norman Sørensen. 2001. "Information aggregation in debate: Who should speak first?" *Journal of Public Economy*, 81(3): 393–421.

Parikh, Rohit, and Paul Krasucki. 1990. "Communication, consensus, and knowledge." *Journal of Economic Theory*, 52(1): 178–189.

Park, Andreas, and Hamid Sabourian. 2011. "Herding and Contrarian Behavior in Financial Markets." *Econometrica*, 79(4): 973–1026.

Rabin, Matthew, and Joel Schrag. 1999. "First Impressions Matter: A Model of Confirmatory Bias." *Quarterly Journal of Economics*, 94: 37–82.

Ritter, Jay R., and Ivo Welch. 2002. "A Review of IPO Activity, Pricing, and Allocations." *Journal of Finance*, 57(4): 1795–1828.

Robin, Eugene D. 1984. *Matters of Life and Death: Risks vs. Benefits of Medical Care.* New York:Freeman and Co.

Rob, Rafael. 1991. "Learning and Capacity Expansion under Demand Uncertainty." *Review of Economic Studies*, 58(4): 655–675.

Romano, Maria Grazia. 2007. "Learning, Cascades and Transaction Costs." *Review of Finance*, 11(3): 527–560.

Rosenberg, Dinah, and Nicolas Vieille. 2019. "On the efficiency of social learning." *Econometrica*, 87(6): 2141–2168.

Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille. 2009. "Informational externalities and emergence of consensus." *Games and Economic Behavior*, 66(2): 979–994.

Ryan, Bryce, and Neal C. Gross. 1943. "The Diffusion of Hybrid Seed Corn in Two Iowa Communities." *Rural Sociology*, 8: 15–24.

Sacerdote, Bruce. 2014. "Experimental and Quasi-Experimental Analysis of Peer Effects: Two Steps Forward?" *Annual Review of Economics*, 6: 253–272.

Salganik, Matthew J., Peter Sheridan Dodds, and Duncan J. Watts. 2006. "Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market." *Science*, 311(5762): 854–856.

Scharfstein, David S., and Jeremy C. Stein. 1990. "Herd behavior and Investment." *American Economic Review*, 80: 465–479.

Sgroi, Daniel. 2002. "Optimizing Information in the Herd: Guinea Pigs, Profits, and Welfare." *Games and Economic Behavior*, 39(1): 137–166.

Sheshinski, Eytan, and Yoram Weiss. 1977. "Inflation and Costs of Price Adjustment." *Review of Economic Studies*, 44(2): 287–303.

Smith, Lones, and Peter Norman Sørensen. 2000. "Pathological Outcomes of Observational Learning." *Econometrica*, 68(2): 371–398.

Smith, Lones, and Peter N. Sørensen. 2020. "Rational Social Learning with Random Sampling." Working paper, University of Wisconsin.

Smith, Lones, Peter Norman Sørensen, and Jianrong Tian. 2021. "Informational herding, optimal experimentation, and contrarianism." *The Review of Economic Studies*, 88(5): 2527–2554.

Song, Yangbo. 2016. "Social learning with endogenous observation." *Journal of Economic Theory*, 166: 324–333.

Sorensen, Alan T. 2006. "Social learning and health plan choice." *RAND Journal of Economics*, 37(4): 929–945.

Spence, Michael. 1978. "Job market signaling." In *Uncertainty in economics.* , ed. Peter Diamond and Michael Rothschild, 281–306. Cambridge, MA:Academic Press.

Sunstein, Cass R. 2009. *A Constitution of Many Minds: Why the Founding Document Doesn't Mean What It Meant Before.* Princeton, NJ:Princeton University Press.

Sunstein, Cass R. 2019*a*. *Conformity: The Power of Social Influences.* Princeton, NJ:NYU Press.

Sunstein, Cass R. 2019*b*. *How change happens.* Cambridge, MA:MIT Press.

Talley, Eric. 1999. "Presidential Cascades: An Appraisal." *Southern California Law Review*, 73: 87–137.

Taylor, Richard. 1979. *Medicine out of Control.* Melbourne:Sun Books.

Tucker, Catherine, Juanjuan Zhang, and Ting Zhu. 2013. "Days on market and home sales." *RAND Journal of Economics*, 44(2): 337–360.

Veeraraghavan, Senthil K., and Laurens G. Debo. 2011. "Herding in Queues with Waiting Costs: Rationality and Regret." *Manufacturing & Service Operations Management*, 13(3): 329–346.

Vermeule, Adrian. 2012. *Law and the Limits of Reason.* . Reprint ed., Oxford University Press.

Vives, Xavier. 1993. "How fast do rational agents learn?" *Review of Economic Studies*, 60: 329–347.

Vives, Xavier. 1997. "Learning from others: a welfare analysis." *Games and Economic Behavior*, 20: 177–200.

Vives, Xavier. 2010. *Information and Learning in Markets: The Impact of Market Microstructure.* Princeton, NJ:Princeton University Press.

Wagner, Peter A. 2018. "Who goes first? Strategic delay under information asymmetry." *Theoretical Economics*, 13(1): 341–375.

Wang, Tao. 2017. "Information revelation through bunching." *Games and Economic Behavior*, 102: 568–582.

Waynforth, David. 2007. "Mate Choice Copying in Humans." *Human Nature*, 18(3): 264–271.

Weizsäcker, Georg. 2010. "Do we follow others when we should? A simple test of rational expectations." *American Economic Review*, 100(5): 2340–2360.

Welch, Ivo. 1992. "Sequential Sales, Learning, and Cascades." *Journal of Finance*, 47(2): 695–732.

Wit, Jorgen. 1997. "Herding behavior in a roll-call voting game." *Department of Economics, University of Amsterdam*.

Witte, Klaudia, Nina Kniel, and Ilka M. Kureck. 2015. "Mate-choice copying: Status quo and where to go." *Current Zoology*, 61(6): 1073–1081.

Wolitzky, Alexander. 2018. "Learning from Others' Outcomes." *American Economic Review*, 108(10): 2763–2801.

Wu, Jiemai. 2021. "Correlated Information Cascades." *Available at SSRN 3951879*.

Young, H. Peyton. 2009. "Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning." *American Economic Review*, 99(5): 1899–1924.

Zhang, Jianbo. 1997. "Strategic Delay and the Onset of Investment Cascades." *RAND Journal of Economics*, 28(1): 188–205.

Zhang, Juanjuan. 2010. "The Sound of Silence: Observational Learning in the U.S. Kidney Market." *Marketing Science*, 29(2): 315–335.

Zhang, Jurui, Yong Liu, and Yubo Chen. 2015. "Social Learning in Networks of Friends versus Strangers." *Marketing Science*, 34(4): 573–589.

Ziegelmeyer, Anthony, Christoph March, and Sebastian Krügel. 2013. "Do We Follow Others When We Should? A Simple Test of Rational Expectations: Comment." *American Economic Review*, 103(6): 2633–2642.

A Online Appendix

A.1 Further discussion of Section 6 on limited communication or observation of past actions

In the next two subsections, we discuss imperfectly rational models, as referred to in § 6.2.3, that make direct heuristic assumptions about the agent's mapping from observed actions and payoffs into the agent's actions. Then, in § A.1.3, we discuss a rational setting with partial communication of past signals.

A.1.1 Heuristic action rules based on observation of averages of past payoffs and of action frequency

If rational agents could observe the average payoffs experienced from possible actions undertaken by many in a large population, they would always choose the right action. However, it is possible that agents fall short of rationality. In the model of Ellison and Fudenberg (1993), there is a continuum of agents, where each period some given fraction of them has the opportunity to update their choices. Agents observe the average payoff of the two actions in the population from the previous period, and their popularities. Agents use rules of thumb for updating their actions, which take the form of specified probability of switching action based on these two variables. The paper explores conditions on the weighting of these considerations that promote or hinder long-run correct decisions.

A.1.2 Heuristic action rules based on sampling of past actions and payoffs

In Ellison and Fudenberg (1995), there is a large number of agents who take simultaneous actions at each discrete period. Each period, some fraction of agents exogenously stick to their current action, and the remainder observes an unbiased sample of the latest actions and payoffs of N and choose action based on this and based on their own latest payoff. Agents follow the following heuristic decision rule. If all the agents in an agent's sample takes the same action, the agent follows that action. If both actions are selected by at least one agent in the sample, the agent chooses whichever action has higher average payoff based on observed reports and the agent's own latest experienced payoff.

When the sample size N is small, learning from the experience of others causes the system to evolve to universal adoption of the correct action. In contrast, when the sample size is large, there is strong pressure toward diversity of behavior.[67] So the system never fixes on the correct action. This analysis is valuable in illustrating how non-obvious interesting conclusions about efficiency derive from reasonably plausible heuristic assumptions. On the other hand, when there is a 'split decision' in an agent's sample, it would be reasonable for an agent to take into account that a preponderance of 99 Adopts to 1 Reject (for example) might suggest almost conclusively that adopt is superior.[68] This would tend to oppose the diversity effect discussed above.

A.1.3 Recent past signals observable

The SBM and the BHW model are based on the premise that agents do not directly communicate their private signals. If the signals of agents are fully communicated to their followers, learning would be efficient, and there would be asymptotic learning.

In many practical contexts, more information is indeed transmitted than just action choices. For example, people are often free to talk about their private information, though owing to time and cognitive constraints, such communication may be imperfect. This raises the question of whether incorrect cascades still occur, and whether there is asymptotic learning.

[67] This is because if one action is very unpopular, with large N many adherents of the more popular action "hear about" the unpopular one and potentially convert to it.

[68] Ellison and Fudenberg consider a related effect which they call 'popularity weighting.'

We consider here the possibility of limited communication of private signals. Every private signal is directly communicated, which in principle could bring about asymptotic learning. However, we consider a case in which each signal is passed on to just one other agent, and provide an example in which incorrect cascades still form and last forever.

Consider the SBM of Section 2, except that in addition to observing all past actions, each agent observes the private signal of the agent's immediate predecessor. In contrast with the SBM, we assume that when indifferent, each agent always follows her own signal. This assumption is convenient but not crucial.

After actions HH or LL, I_3 infers that the first two signals were hh or $\ell\ell$ respectively, and falls into a cascade which could easily be incorrect. This cascade can be broken. To see this, suppose that the first two actions are HH, and that I_3 observes ℓ. So I_3 chooses H. If I_4 observes ℓ, this combines with observation of I_3's ℓ to make I_4 indifferent, so I_4 rejects, breaking the cascade.

Consider next the case in which I_3 observes h as well. Now, I_4 adopts, because I_4 knows that the first 3 signals were all h. At this point the cascade of adoption continues forever. I_5 knows that the first two signals were hh, and that the third and fourth signals were not $\ell\ell$. (I.e., they were either ℓh, hh, or $h\ell$.) So the net evidence in favor of state H is at least slightly stronger than two h signals. So even in the worst case in which I_5 directly observes two ℓ signals (I_5's and I_4's signals), I_5 strictly prefers H. Similar reasoning shows that all subsequent agents adopt.

The general point that inefficient cascades can form and last forever does not require this arbitrary tie-breaking convention.[69] Intuitively, the ability to observe the predecessor's signal is a double-edged blade. The immediate effect is to increase an

[69]The following modification illustrates the point in a setting with no ties. Suppose that agents differ very slightly in their precisions, and that agents act in inverse order of their precisions, from low to high. Now agents are never indifferent. Agents who are close to indifference strictly prefer to follow their own signals, and the same analysis holds.

agent's information, which increases the probability that an agent decides correctly. But to the extent that this is the case, it later makes the action history more informative. That eventually encourages later agents to fall into a cascade, which blocks asymptotic learning.

In particular, the action sequence becomes informative enough that agents fall into cascades despite their access to an extra signal from the predecessor. The actions history eventually overwhelms an agent's information (even inclusive of observation of a predecessor's signal), at which point the accuracy of the social belief stops growing.

Another way to think of this is that being able to observe predecessors' signals is somewhat like being able to observe multiple private signals instead of one. This does not fundamentally change the argument for why incorrect cascades form—that the accumulation of information implicit in past actions must eventually overwhelm the signal(s) of a single agent.