

Erik Eyster, Andrea Galeotti, Navin Kartik and Matthew Rabin

Congested observational learning

**Article (Accepted version)
(Refereed)**

Original citation:

Eyster, Erik, Galeotti, Andrea, Kartik, Navin and Rabin, Matthew (2014) Congested observational learning. [Games and Economic Behavior](#) . ISSN 0899-8256 (In Press)

© 2014 The Authors

This version available at: <http://eprints.lse.ac.uk/58748/>

Available in LSE Research Online: August 2014

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

This document is the author's final accepted version of the journal article. There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

Congested Observational Learning

Erik Eyster, Andrea Galeotti, Navin Kartik, and Matthew Rabin*

June 25, 2014

Abstract

We study observational learning in environments with congestion costs: an agent's payoff from choosing an action decreases as more predecessors choose that action. Herds cannot occur if congestion on every action can get so large that an agent prefers a different action regardless of his beliefs about the state. To the extent that switching away from the more popular action reveals private information, it improves learning. The absence of herding does not guarantee complete (asymptotic) learning, however, as information cascades can occur through perpetual but uninformative switching between actions. We provide conditions on congestion costs that guarantee complete learning and conditions that guarantee bounded learning. Learning can be virtually complete even if each agent has only an infinitesimal effect on congestion costs. We apply our results to markets where congestion costs arise through responsive pricing and to queuing problems where agents dislike waiting for service.

*Eyster: Department of Economics, London School of Economics, Houghton Street, London WC2A 2AE United Kingdom (email: e.eyster@lse.ac.uk); Galeotti: Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3S, United Kingdom (email: agaleo@essex.ac.uk); Kartik: Department of Economics, Columbia University, 420 W. 118th Street, New York, NY 10027, USA (email: nkartik@columbia.edu); Rabin: Department of Economics, University of California, Berkeley, 549 Evans Hall #3880, Berkeley, CA 94720-3880 USA, (email: rabin@econ.berkeley.edu). We thank Wouter Dessen for his discussion at the 2012 Utah Winter Business Economics Conference. We also received helpful comments from Tristan Gagnon-Bartsch; audiences at the LSE, University of Essex, and University of Vienna; and an anonymous Associate Editor and three referees. Teck Yong Tan provided exemplary research assistance. Andrea Galeotti is grateful to the European Research Council for support through ERC-starting grant (award no. 283454) and to the Leverhulme Trust for support through the Philip Leverhulme Prize.

1 Introduction

We examine how rational agents learn from observing the actions—but not directly the information—of other rational agents. The focus of our study is a class of payoff interdependence: as more of an agent’s predecessors choose one action, the agent’s payoff from choosing that action decreases. We term this kind of payoff interdependence *congestion costs*. They naturally arise as direct economic costs in many contexts. For example, as more individuals purchase a product, prices may increase (Avery and Zemsky, 1998), short-term supplies may run out, quality of service may worsen, or waiting times may lengthen. There are empirical estimates that “queuing costs” have significant effects on consumer behavior (Aksin et al., 2013; Lu et al., 2013). Another example is in the context of market entry: when firms choose whether to enter a new market, expected profits are likely to be decreasing in the number of other entrants. Congestion costs can also reflect a pure taste for “anti-conformity”, as agents sometimes have intrinsic preferences for avoiding options to which others flock.

Our model in Section 2 builds on the canonical models of observational learning (Banerjee, 1992; Bikhchandani et al., 1992; Smith and Sørensen, 2000). A sequence of agents each choose between two actions, A and B . One of the actions is superior to the other; all agents share common preferences on this dimension, but each agent has imperfect private information about which action is superior. Each agent acts based on his private signal and the observed choices of all predecessors. We assume that private signals have bounded informativeness because this case turns out to be more interesting, but otherwise make no assumptions about their distribution.

We enrich the standard model by assuming that while all agents prefer the superior action *ceteris paribus*, agents may also dislike taking an action more when more of their predecessors have chosen it.¹ We parameterize how much agents care about these congestion costs relative to taking the superior action by a marginal-rate-of-substitution parameter $k \geq 0$. For example, the congestion cost associated with an action may equal k times the number of predecessors who chose that action. When $k = 0$, the model obviously collapses to the standard one without congestion costs. Our primary focus is on the long-run outcomes of observational learning when $k > 0$. Does society eventually learn which action is superior, and how does the presence of congestion affect the long-run frequency of actions, in particular the fraction of agents who choose the superior action?

Section 3 develops some preliminaries about individual decision-making as a function of the private signal and the public history of actions. We turn to the asymptotic properties of observational learning in Section 4. In the canonical model without congestion, all agents eventually take the same action, i.e. herds necessarily form in finite time.² Moreover, learning is bounded in the sense that society never learns with certainty, even asymptotically, which

¹We discuss positive payoff interdependence in the conclusion, Section 7.

²Throughout this introduction, we ease exposition by suppressing technical details such as “almost sure” caveats, and often referring to just “learning” in lieu of “asymptotic learning”.

action is superior. Both these conclusions also hold in the current environment so long as the net congestion cost (i.e. the difference in congestion cost incurred by taking one action rather than the other) that an agent may face is small enough.³

By contrast, if the net congestion cost can become sufficiently high, then it is clear that herds cannot form: should a long-enough sequence of agents take action A , then eventually someone will take action B even if extremely confident that A is superior. It is noteworthy, however, that the impossibility of herding does not imply that society eventually learns which action is superior. What is crucial is whether agents “switch” from one action to the other purely in response to congestion, or also in response to their private information. It is possible that after some time, agents perpetually cycle between the two actions without conveying any information about their private signals. In other words, an information cascade may begin wherein every agent’s action is preordained despite the absence of a herd; this phenomenon can occur no matter how much agents care about congestion relative to taking the superior action, i.e. no matter the value of the marginal-rate-of-substitution parameter $k > 0$. Indeed, there are interesting classes of congestion costs where net congestion can grow arbitrarily large and yet, no matter the value of k , an information cascade with cyclical behavior will necessarily arise in finite time.

Whether such outcomes can occur turns on the incremental effect that any one agent’s action has on the net congestion cost faced by his successors. We identify general properties of congestion costs that ensure bounded learning even when herds cannot arise. Conversely, we provide conditions that guarantee complete learning. What drives complete learning when it arises in our model is not that *every* agent behaves informatively; rather it is an inevitable return to *some* agent behaving informatively. Put differently, even when there is complete learning, it will often be the case that on any sample path of play there are (many) phases of “temporary information cascades”.

A natural question is what happens when any given agent cares little about congestion, i.e. the marginal-rate-of-substitution parameter $k > 0$ vanishes. When net congestion costs are bounded there is bounded learning once k is small enough. However, if net congestion can grow arbitrarily large when sufficiently many consecutive agents play the same action, then there is essentially complete learning as $k \rightarrow 0$. Intuitively, only at near-certain public beliefs can vanishingly small incremental congestion costs produce the sort of uninformative cycling that stalls learning.

There is a sense in which this result can be interpreted as a fragility of the conventional bounded-learning result. [Section 4](#) clarifies precisely how our model is (and is not) continuous at the limit as $k \rightarrow 0$. Here, we illustrate a substantive economic point using the example where the congestion cost of taking an action equals the number of predecessors who have chosen it. In this example, when $k = 0.1$ someone would prefer to patronize a restaurant known to be inferior in order to reduce his “queue” by ten people. When $k = 0.01$ he would

³More precisely, this requires that the net congestion cost be bounded above in absolute value by some threshold that is sufficiently small relative to the marginal-rate-of-substitution parameter k .

only do so in order to reduce his queue by one hundred people. One might expect that insights from the standard model would provide a good approximation for such small k . With regards to asymptotic learning, however, this intuition can be dramatically incorrect. Consider, for instance, the canonical binary-signal structure where each agent receives a signal that matches the true state (i.e. the superior action) $2/3$ of the time. When $k = 0$, the asymptotic public belief that an action is the superior action cannot exceed 0.8 nor drop below 0.2. When $k = 0.1$, on the other hand, this public belief must settle either above 0.96 or below 0.04; when $k = 0.01$, it must settle above 0.996 or below 0.004. Thus, when individual preferences more closely resemble the standard no-congestion-cost model, not only are long-run beliefs more confident, but those beliefs are also further away from the standard no-congestion-cost model. It should be noted, however, that the force whereby smaller k leads to more confident long-run beliefs goes hand in hand with slower learning, because it takes longer for enough congestion to accumulate before agents switch their actions. While our focus on asymptotic learning rather than speed of convergence follows most of the literature on observational learning (see [Lobel et al. \(2009\)](#) for an exception), the focus may be a greater weakness when considering comparative statics in k because of the aforementioned tension.

In [Section 5](#), we discuss three applications that fit into our general framework. The first application interprets congestion costs as being induced by a tax-and-redistribution scheme used by a social planner in the standard observational learning model without congestion. By establishing that agents will asymptotically choose the superior action under certain forms of congestion costs, we propose a simple transfer scheme that a planner can use to obtain the first-best outcome in the long run.⁴ The second application views congestion costs as induced by the evolution of market prices over time. Our analysis here accommodates a class of reduced-form price-setting rules that correspond to a range of market-competition assumptions from monopoly at one end to Bertrand competition at the other. We illustrate how different price-setting rules yield different conclusions about whether prices eventually reveal all agents' information. The third application explores a queuing model where agents are served in sequence, but service only occurs with some probability in each period. Congestion costs here arise from agents' dislike for delay in being served. We illustrate how different assumptions about the observability of predecessors being served yield different conclusions about asymptotic learning.

A limitation of our baseline model is the assumption that agents must necessarily choose between options A and B regardless of how large the congestion cost of either option is. In other words, it effectively assumes that all other options for an agent are dominated, no matter how severe the congestion is on the two options with uncertain payoffs. In [Section 6](#), we explore the robustness of our learning results to the presence of an outside option that provides some fixed utility level, independent of both the state of the world and the behavior of other agents. We make two assumptions, both of which we view as sensible for many applications: (i) congestion

⁴See [Naor \(1969\)](#) for early work examining optimal taxation (“tolls”) to reduce the negative externality caused by congestion in a queuing model without learning.

on A and B will eventually decay fully if agents avoid that option altogether; (ii) if there is no congestion on either A or B , the outside option is dominated. Our main results carry over to this setting; in particular, we prove that when net congestion costs are unbounded, even though agents may often choose the outside option—“balk”, in the terminology of queuing theory—there is essentially complete learning as $k \rightarrow 0$, just like in the baseline model without an outside option.

Related Literature. There are few prior studies of observational learning with direct congestion or queuing costs. Gaigl (2009) assumes that congestion costs take a particular functional form that is subsumed by our general formulation; specifically, he analyzes what we call the *linear absolute-cost* example (Example 2 in Subsection 2.1), where the congestion cost of an action is a linear function of the number of predecessors who have taken that action.⁵ For continuous signal structures, he discusses when information cascades and herds can occur but does not address asymptotic learning; for binary signals, he also provides results on learning. Besides accommodating richer signal structures, our analysis reveals that the nature of congestion costs is key: the linear absolute-cost example satisfies two properties—congestion is *unbounded* and has *gaps* (defined in Section 4)—that need not hold more generally but matter crucially for conclusions about learning, information cascades, and herds. For instance, complete learning does not arise in the linear absolute-cost example, but does for other congestion costs. A general analysis of different kinds of congestion costs requires distinct techniques, yields broader theoretical insights, and permits us to apply our results to different economic applications.

Veeraraghavan and Debo (2011) and Debo et al. (2012) develop queuing-with-learning models where agents observe only an aggregate statistic of predecessors’ choices but not the entire history. Bayesian inference in this setting is extremely complex, and hence these papers do not analyse asymptotic learning but instead characterise some properties of equilibrium play in early rounds. Cripps and Thomas (2014) characterize equilibria in a queuing problem with strategic experimentation; in their setting, agents decide both whether to join the queue and whether (and when) to quit it before being served. More broadly, Hassin and Haviv (2003) provide an introduction to strategic queuing models, albeit generally without learning.

Our work also relates to Avery and Zemsky (1998), who build on Glosten and Milgrom’s (1985) model of sequential trade for an asset of common but unknown value. In Avery and Zemsky’s introductory example, a market-maker sets the price of a risky asset at the start of every period to equal its expected value based on all prior trades. As the price fully incorporates all public information, each trader acts informatively, buying when his private information is favourable and selling when it is unfavourable.⁶ The market price thus converges to the asset’s true value. Such market prices play a similar role to congestion costs in our model: holding

⁵We learned of Gaigl’s work only after circulating a prior draft of this paper.

⁶Thus, in this example, the market-maker loses money on average. Avery and Zemsky’s (1998) richer model with noise traders and bid-ask spreads does not share this feature. The authors’ focus is on how multidimensional private information allows for herding even with informationally-efficient prices. Park and Sabourian (2011) clarify the information structures needed for such results.

fixed a trader’s belief about the asset’s value, buying the asset becomes less desirable as more predecessors buy the asset. Our model can be seen as extending this introductory example from [Avery and Zemsky \(1998\)](#) beyond the realm of markets and specific theories of price formation. Doing so, we show on the one hand that complete learning can obtain even in settings where most agents act uninformatively, and on the other hand that different mechanisms of price formation can substantially alter the conclusion of complete learning.

There are models of observational learning without congestion costs in which complete learning obtains. [Lee \(1993\)](#) derives such a result when agents’ action spaces are a continuum—rich enough to reveal their posteriors—and preferences satisfy some reasonable properties. Even with only a finite number of actions, [Smith and Sørensen \(2000\)](#) and [Goeree et al. \(2006\)](#) respectively show that complete learning obtains when private beliefs are unbounded (and agents have the same preferences) or agents’ preferences include a full-support private-values component (while private beliefs remain bounded).⁷ Congestion costs in our framework generate heterogeneity in agents’ preferences, but do so endogenously and in a sample-path-dependent manner. It is worth emphasizing that large—even unbounded—total congestion costs do not imply complete learning; by creating a “wedge” between agents’ utilities from the two actions, large congestion costs are compatible with information cascades (but not herds).⁸ Furthermore, in the settings explored by [Lee \(1993\)](#), [Smith and Sørensen \(2000\)](#), and [Goeree et al. \(2006\)](#), *every* agent’s action reveals some information about his private signal; as already mentioned, this is typically not the case here even when complete learning obtains.

Finally, this paper contributes to a growing theoretical literature on observational learning when there is direct payoff interdependence between agents. A significant fraction of this literature has focussed on sequential elections (e.g. [Dekel and Piccione, 2000](#); [Callander, 2007](#); [Ali and Kartik, 2012](#), and the references therein), but other work also studies coordination problems ([Dasgupta, 2000](#)), common-value auctions ([Neeman and Orosel, 1999](#)), settings with network externalities ([Choi, 1997](#)), and when agents partially internalize the welfare of future agents ([Smith and Sørensen, 2008](#)). Congestion-cost models such as ours focus on a different kind of payoff interdependence and on environments where agents only care about past actions. While the latter is a limitation for some applications, it is appropriate in other contexts and permits a fairly general treatment of the payoff interdependence we study.⁹

⁷Intuitively, information cascades cannot occur under unbounded private beliefs or full-support preference shocks because whatever the current public belief, an agent can receive either a sufficiently strong signal or preference shock to overturn it. [Smith and Sørensen \(2000\)](#) and [Herrera and Hörner \(2012\)](#) observe that, under bounded private beliefs and common preferences, the absence of information cascades is compatible with a failure of complete learning. [Acemoglu et al. \(2011\)](#) show that complete learning can occur under bounded private beliefs and common preferences if not all agents necessarily observe all predecessors’ choices.

⁸That information cascades can arise in the absence of herds due to preference heterogeneity has been noted in other settings of observational learning, e.g. by [Smith and Sørensen \(2000\)](#), [Cipriani and Guarino \(2008\)](#) and [Dasgupta and Prat \(2008\)](#).

⁹On the experimental front, [Drehmann et al. \(2007\)](#) conduct an internet experiment that includes a treatment with congestion costs, which they show reduce the average length of runs of consecutive actions modestly compared to the no-congestion benchmark. They also include treatments with forward-looking payoff externalities, in which they find that subjects’ behavior do not differ significantly from myopic behavior. Accordingly, they suggest that

2 The Model

A payoff-relevant state of the world, $\theta \in \{-1, 1\}$, is (without loss of generality) drawn from a uniform prior. A countable infinity of agents, $i = 1, 2, \dots$, take actions in sequence, each observing the entire history of actions. Before choosing his action, each agent i gets a private signal that is independently and identically distributed conditional on the state. Following [Smith and Sørensen \(2000\)](#), we work directly with the random variable of *private beliefs*, which is an agent’s belief that $\theta = 1$ after observing his signal but ignoring the history of play, computed by Bayes rule using the uniform prior and the signal distributions. Denote the private belief of agent i as $b_i \in (0, 1)$. Given the state $\theta \in \{-1, 1\}$, the private-belief stochastic process $\langle b_i \rangle$ is conditionally i.i.d. with conditional c.d.f. $F(\cdot|\theta)$. We assume that no private signal perfectly reveals the state of the world, which implies that $F(\cdot|1)$ and $F(\cdot|-1)$ are mutually absolutely continuous and have common support. Denote the support’s convex hull by $[\underline{b}, \bar{b}] \subseteq [0, 1]$. To avoid trivialities, signals must be informative, which implies that $\underline{b} < 1/2 < \bar{b}$. To focus on the most interesting case, we assume *bounded private beliefs*: $\underline{b} > 0$ and $\bar{b} < 1$; the case of unbounded private beliefs ($\underline{b} = 0$ and $\bar{b} = 1$) is discussed in the conclusion. Notice that this setting allows for continuous or discrete signals.

We assume for now that each agent faces a binary choice; [Section 6](#) discusses a generalization. Denote each agent’s action by $a_i \in \{-1, 1\}$ and let $a^i := (a_1, \dots, a_i)$ denote a history. Agent i ’s preferences are given by a von-Neumann-Morgenstern utility function

$$u_i(a^i, \theta) := \mathbb{1}_{\{a_i = \theta\}} - kc(a^i),$$

where $\mathbb{1}_{\{\cdot\}}$ denotes the indicator function, $c(\cdot)$ is a state-independent congestion cost function, and $k > 0$ is a scalar parameter. Gross of congestion costs, the gain from taking the *superior* action (i.e. the action that matches the state) is normalized to one. The assumption that $c(\cdot)$ depends only upon a^i implies that congestion is “backward looking” in the sense of only depending on predecessors’ choices. Note that because the domain of $c(\cdot)$ varies with an agent’s index, different agents may be affected differently in terms of congestion by common predecessors, and furthermore, different agents may trade off the gain from taking the superior action relative to congestion differently. The standard model without congestion obtains when $k = 0$. For a fixed game, the scalar k could be folded into the cost function $c(\cdot)$, but our parametrization allows us to discuss a sequence of congestion games converging to a no-congestion game by holding $c(\cdot)$ fixed and letting $k \rightarrow 0$.

Insofar as congestion is concerned, agent $i + 1$ ’s choice depends only on the *net congestion cost* he faces, i.e. the additional cost of choosing $a = 1$ rather than $a = -1$, which is given by

$$\Delta(a^i) := c((a^i, 1)) - c((a^i, -1)).$$

a purely backward-looking analysis might make reasonably good predictions even in settings with forward-looking considerations. [Owens \(2010\)](#) presents a laboratory experiment on observational learning with payoff externalities, finding that decisions are highly responsive to both positive and negative payoff externalities.

We capture the negative externality from congestion by assuming that an extra action causes the net congestion cost of taking that action to weakly rise. Formally:

Assumption 1 (Monotonicity). *For all a^i , $\Delta((a^i, 1)) \geq \Delta(a^i) \geq \Delta((a^i, -1))$.*

Our other maintained assumption is that the marginal impact of any individual's choice on net congestion cost is bounded above. Formally:

Assumption 2 (Bounded increments). $\sup_{a^i} \max\{\Delta((a^i, 1)) - \Delta(a^i), \Delta(a^i) - \Delta((a^i, -1))\} < \infty$.

2.1 Leading examples

Below, we introduce two leading examples that satisfy all the maintained assumptions. These examples are primarily meant to help illustrate the conditions and results that will follow; we later discuss variations and other specifications. In both of these examples, agents care only about the frequency with which their predecessors have chosen one action over the other (rather than which predecessors chose which action), but the examples differ in whether frequency is measured in proportional or absolute terms.

Example 1. In the *linear proportional-cost model*, $c(a^1) = 0$ and for each $i \geq 2$,

$$c(a^i) = \frac{\sum_{j=1}^{i-1} \mathbb{1}_{\{a_j = a_i\}}}{i-1}.$$

If we denote $\rho(a^i)$ as the proportion of agents $1, \dots, i$ who have chosen $a = 1$ under a^i , then for each $i \geq 2$, $\Delta(a^i) = 2\rho(a^i) - 1$. More generally, instead of $\Delta(a^i)$ being linear in $\rho(a^i)$, we could have $\Delta(a^i) = f(\rho(a^i))$ for some function $f : (0, 1) \rightarrow \mathbb{R}$ that is strictly increasing but otherwise arbitrary; this defines the general *proportional-costs model*.¹⁰ ||

Example 2. In the *linear absolute-cost model*, let $\eta(a^i) = \sum_{j=1}^i \mathbb{1}_{\{a_j = 1\}}$ be the number of agents who have chosen action 1 in the sequence a^i , and define $c(a^1) = 0$ and for each $i \geq 2$,

$$c(a^i) = \begin{cases} \eta(a^{i-1}) & \text{if } a_i = 1 \\ i - 1 - \eta(a^{i-1}) & \text{if } a_i = -1. \end{cases}$$

Here, congestion depends upon the number of agents who have chosen each action rather than the fraction, so that net congestion $\Delta(a^i) = 2\eta(a^i) - i$. A general *absolute-costs model* has $\Delta(a^i) = f(\eta(a^i)) - f(i - \eta(a^i))$ for some strictly increasing $f : \mathbb{R}_+ \rightarrow \mathbb{R}$. ||

¹⁰Note that $f(\cdot)$ is only defined here on the interior of the unit interval. To avoid some inessential complications, in this example we exogenously set the first two agents' choices to be $a_1 = 1$ and $a_2 = -1$.

3 Individual Decision-Making

Agent i 's decision depends upon his beliefs about which action is superior as well as upon net congestion costs. Let b be agent i 's private belief that the state is $\theta = 1$, which depends on i 's private signal alone; let q be the public belief that the state is $\theta = 1$, which depends on the inference that agent i makes from the history a^{i-1} . Furthermore, let $l = (1 - q)/q$ be the *public likelihood ratio*, which is the inverse of the relative likelihood of state 1; low l means it is more likely that $\theta = 1$. By Bayesian updating, i 's posterior belief that the state is $\theta = 1$ is given by

$$p(b, l) = \frac{b}{b + (1 - b)l}. \quad (1)$$

As is intuitive, $p(b, l)$ is strictly increasing in b and strictly decreasing in l . The following lemma describes how posterior beliefs determine action choice.¹¹

Lemma 1. *An agent who has private belief b , public likelihood ratio l , and net congestion cost Δ chooses action 1 if and only if $p(b, l) \geq p^*(\Delta; k) := 1/2 + k\Delta/2$.*

Proof. It suffices to compute

$$\begin{aligned} \mathbb{E}[u_i(a_i = 1) - u_i(a_i = -1)] &= p(b, l) [1 - kc(a^{i-1}, 1) - (0 - kc(a^{i-1}, -1))] \\ &\quad + (1 - p(b, l)) [0 - kc(a^{i-1}, 1) - (1 - kc(a^{i-1}, -1))] \\ &= 2p(b, l) - (1 + k\Delta), \end{aligned}$$

which immediately implies the result. *Q.E.D.*

Using Lemma 1, we can derive the net congestion cost that renders an agent indifferent between action 1 and action -1 given a public likelihood ratio l and the private belief most favorable to action 1. Formally, for any $l \in \mathbb{R}_+$, define $\bar{\Delta}(l; k)$ to be the unique solution to $p(\bar{b}, l) = p^*(\bar{\Delta}(l; k); k)$.¹² Given a net congestion cost of $\bar{\Delta}(l; k)$ and public likelihood l , an agent is indifferent between choosing $a = -1$ and $a = 1$ when he receives the private belief most favorable to action 1. For any b, l , and $\Delta \geq \bar{\Delta}(l; k)$, we have $p(b, l) \leq p^*(\Delta; k)$, which, by Lemma 1, implies that the agent will choose $a = -1$ regardless of his private belief.

Similarly, for any $l \in \mathbb{R}_+$, define $\underline{\Delta}(l; k)$ to be the unique solution to $p(\underline{b}, l) = p^*(\underline{\Delta}(l; k); k)$, which is well defined because $\underline{b} > 0$. In this case, $\underline{\Delta}(l; k)$ is the net congestion cost that renders an agent indifferent between action 1 and action -1 given likelihood ratio l and the private belief most favourable to action -1 . Hence, for any b, l , and $\Delta \leq \underline{\Delta}(l; k)$, an agent will choose $a = 1$ regardless of his private belief.

¹¹Throughout the paper, we resolve agents' indifference over actions in the manner most convenient for exposition. Note that the choice of how to resolve indifference is irrelevant with continuous signals; with discrete signals, it is generically irrelevant.

¹²Uniqueness is guaranteed because $p^*(\cdot; k)$ is strictly increasing and unbounded above and below.

We call an agent's action *informative* if it depends non-trivially upon his private beliefs, or equivalently, if $\Delta \in (\underline{\Delta}(l; k), \overline{\Delta}(l; k))$. The next lemma provides a number of important properties of the net congestion threshold functions.

Lemma 2. *The net congestion threshold functions, $\overline{\Delta}(\cdot; \cdot)$ and $\underline{\Delta}(\cdot; \cdot)$, satisfy the following:*

1. *For any $l > 0$ and $k > 0$, $\overline{\Delta}(l; k) > \underline{\Delta}(l; k)$. An agent's action choice is informative if and only if $\Delta \in (\underline{\Delta}(l; k), \overline{\Delta}(l; k))$; if $\Delta \geq \overline{\Delta}(l; k)$, he chooses $a = -1$ for any private belief; if $\Delta \leq \underline{\Delta}(l; k)$, he chooses $a = 1$ for any private belief.*
2. *For any $k > 0$: (i) $\overline{\Delta}(0; k) = \underline{\Delta}(0; k) = 1/k$; (ii) as functions of l , both $\overline{\Delta}(l; k)$ and $\underline{\Delta}(l; k)$ are continuous, strictly decreasing, and converge to $-1/k$ as $l \rightarrow \infty$; and (iii) $\overline{\Delta}(l; k) - \underline{\Delta}(l; k) \rightarrow 0$ as $l \rightarrow \infty$ or as $l \rightarrow 0$ and is quasi-concave in l .*
3. *For any $l > 0$, $\overline{\Delta}(l; k) - \underline{\Delta}(l; k) \rightarrow \infty$ as $k \rightarrow 0$.*
4. *$\overline{\Delta}(l; \cdot)$ is strictly decreasing (resp. increasing) for any $l > 0$ that is strictly smaller (resp. larger) than $\frac{\bar{b}}{1-\bar{b}}$; similarly $\underline{\Delta}(l; \cdot)$ is strictly decreasing (resp. increasing) for any $l > 0$ that is strictly smaller (resp. larger) than $\frac{\underline{b}}{1-\underline{b}}$.*

Proof. The definitions of $\overline{\Delta}(\cdot)$ and $\underline{\Delta}(\cdot)$ imply

$$1 + k\overline{\Delta}(l; k) = 2p(\bar{b}, l) \quad \text{and} \quad 1 + k\underline{\Delta}(l; k) = 2p(\underline{b}, l), \quad (2)$$

into which we can substitute [Equation 1](#) to derive

$$\overline{\Delta}(l; k) = \frac{\bar{b} - (1 - \bar{b})l}{k(\bar{b} + (1 - \bar{b})l)} \quad \text{and} \quad \underline{\Delta}(l; k) = \frac{\underline{b} - (1 - \underline{b})l}{k(\underline{b} + (1 - \underline{b})l)}.$$

The lemma's first two and fourth parts are straightforward to verify using these formulae and [Lemma 1](#). We next combine the two equations in [\(2\)](#) to derive

$$k(\overline{\Delta}(l; k) - \underline{\Delta}(l; k)) = 2(p(\bar{b}, l) - p(\underline{b}, l)).$$

The third part of the lemma follows from the observation that the right hand side above is strictly positive. *Q.E.D.*

Figure [1\(a\)](#) illustrates [Lemma 2](#)'s first two observations. For net congestion costs above $\overline{\Delta}(l; k)$ (the bold line), agents choose action -1 regardless of their private belief; for net congestion costs below $\underline{\Delta}(l; k)$ (the dotted line), agents choose action 1 regardless of their private belief; for net congestion costs that lie between the two lines, agents choose actions that depend upon their private beliefs. Figure [1\(b\)](#) illustrates how a change in k affects both $\underline{\Delta}(l; k)$ and $\overline{\Delta}(l; k)$, as described in parts 3 and 4 of [Lemma 2](#). For a given l , when k decreases, $\overline{\Delta}(l; k)$ rotates around $\bar{b}/(1 - \bar{b})$: it increases when the likelihood ratio is below $\bar{b}/(1 - \bar{b})$ and it decreases in the complementary region. The change in $\underline{\Delta}(l; k)$ is analogous. Furthermore, for any $l > 0$, a decrease in k increases the difference between the two thresholds; this difference gets arbitrarily large as k vanishes.

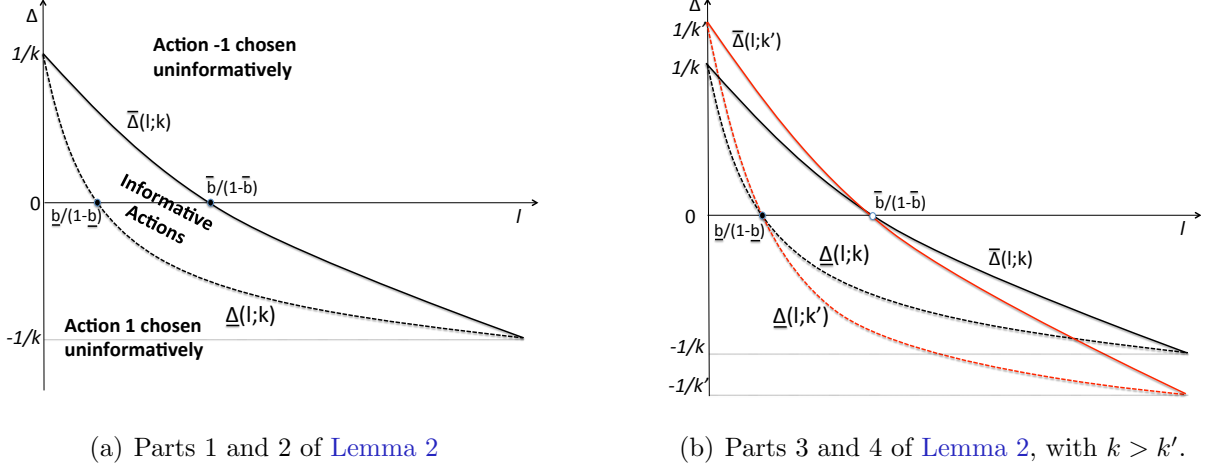


Figure 1 – Net congestion threshold functions.

4 Asymptotic Learning

4.1 Concepts

Without loss of generality, assume the true state is $\theta = 1$. By standard arguments, the public-likelihood-ratio stochastic process, $\langle l_i^k \rangle$, is a (conditional) martingale. Thus, it almost surely converges to a random variable l_∞^k such that $\text{Supp}[l_\infty^k] \subseteq [0, \infty)$.

We say there is *complete learning* if beliefs almost surely converge to the truth, i.e. $\Pr(l_\infty^k = 0) = 1$. We say there is *bounded learning* if there exists $\varepsilon > 0$ such that $\Pr(l_\infty^k > \varepsilon) = 1$, because then beliefs almost surely are bounded away from the truth. There is a *herd* on a sample path if after some (finite) time, all subsequent agents choose the same action. There is an *information cascade* on a sample path if after some (finite) time, no agent's action is informative.

In the benchmark no-congestion model of observational learning, the existing literature has established the following results:

Remark 1. Assume $k = 0$. Since private beliefs are bounded, there is bounded learning. A herd occurs almost surely; moreover, with positive probability, the herd forms on the inferior action. Whether an information cascade can arise depends on the distributions of private beliefs.¹³ ||

¹³Regarding learning and herds, everything stated in Remark 1 except bounded learning follows from Theorem 1 in Smith and Sørensen (2000). To see that learning is bounded, define $\underline{l} \in (0, 1)$ by $p(\underline{b}, \underline{l}) = 1/2$. Note that on any sample path, if $l_i^0 < \underline{l}$ for some i , then $a_i = 1$ independent of i 's private belief, hence $l_{i+1}^0 = l_i^0$. Now define $\hat{l} := \frac{1-p(\underline{b}, \underline{l})}{p(\underline{b}, \underline{l})} \underline{l}$, i.e. \hat{l} is the posterior likelihood ratio obtained from a public likelihood ratio of \underline{l} and the most favorable private belief; observe that $\hat{l} \in (0, \underline{l})$. It follows that there is no sample path in which for some i , $l_i^0 < \hat{l}$. For a characterization of when information cascades can and cannot arise under bounded private beliefs, see Herrera and Hörner (2012).

Note that in the standard model without congestion, an information cascade implies a herd. As we will see, this need not be the case in the current setting.

Aside from studying asymptotic learning for a given $k > 0$, we are also interested in what happens in the no-congestion limit, i.e. as $k \rightarrow 0$. Say that there is *complete learning in the no-congestion limit* if $\lim_{k \rightarrow 0} \Pr(l_\infty^k = 0) = 1$. Say that there is *learning with high probability in the no-congestion limit* if for all $\varepsilon > 0$, $\lim_{k \rightarrow 0} \Pr[l_\infty^k < \varepsilon] = 1$. Finally, there is *bounded learning in the no-congestion limit* if there exists $\varepsilon > 0$ such that $\lim_{k \rightarrow 0} \Pr(l_\infty^k > \varepsilon) = 1$.

Complete learning in the no-congestion limit captures the notion that limit learning occurs almost surely. In particular, if there is complete learning for all (small enough) $k \geq 0$, then there is complete learning in the no-congestion limit. Learning with high probability in the no-congestion limit captures the weaker notion that the sequence of random variables l_∞^k converges in probability to 0 as $k \rightarrow 0$. Notice that even this weaker notion of learning represents a discontinuity with the standard model under bounded beliefs, where the asymptotic public belief is bounded away from the truth (see [Remark 1](#)). Finally, a case of bounded learning in the no-congestion limit resembles the standard model without congestion.

4.2 An overview of the results

Before turning to the formal analysis, this subsection discusses the main intuitions for how different forms of congestion costs affect learning. In conveying intuitively the ideas behind the theorems of Sections 4.3–4.5, we deliberately ignore some subtleties in this subsection.

For a given $k > 0$, there are two potential reasons why complete learning may not occur. The first is standard: it is possible that agents will eventually herd uninformatively. Whether a herd can occur in the present context depends on the cumulative effect of agents' actions on net congestion costs. During a putative herd on action 1, say, agents' net congestion cost is increasing, i.e., it becomes increasingly more costly to choose action 1 relative to -1 . If the net congestion cost never becomes prohibitively high relative to the public likelihood ratio (which remains unchanged while agents behave uninformatively), the herd persists. [Figure 2\(a\)](#) illustrates this possibility: the solid circles with arrows depict the evolution of (l, Δ) , i.e. the public likelihood ratio and net congestion, during a herd on action 1.

Clearly, if the net congestion cost would eventually become sufficiently high (resp. sufficiently low) should enough consecutive agents choose action 1 (resp. action -1), then eventually some agent would break a putative herd on action 1 (resp. -1).¹⁴

The impossibility of uninformative herding is not sufficient for complete learning, however. The second reason why learning can stop is that agents may perpetually switch between the two actions while never conveying any information about private beliefs. [Figure 2\(b\)](#) illustrates this possibility in a particularly stark fashion: eventually each agent just takes the opposite action of

¹⁴More precisely, there cannot be a herd on action 1 following the history a^i if $\lim \Delta((a^i, 1, 1, \dots)) > 1/k$, and hence there can never be a herd on action 1 if $\inf_{a^i} \lim \Delta((a^i, 1, 1, \dots)) > 1/k$. Analogous statements hold for herds on action -1 .

his predecessor, as depicted under parameter value k by the oscillation between the two hollow circles. Such an outcome is only possible for certain net congestion cost functions; [Theorem 3](#) below identifies a sufficient condition. Notice that for any $k > 0$, there is a gap between the $\bar{\Delta}$ and $\underline{\Delta}$ curves at any interior public likelihood ratio. Thus, perpetual uninformative oscillation at any public likelihood ratio would be precluded for all $k > 0$ if the incremental effect of any agent’s action on net congestion eventually became negligible, because this would ensure that net congestion cannot perpetually “jump” across the thresholds. [Theorem 2](#) below develops this idea into a result on complete learning for all $k > 0$.

Even when uninformative oscillations between actions can occur at some public likelihood ratio, say l , for a given value of k , they can only do so when (some) agents have incremental effects on net congestion cost that is larger than $\bar{\Delta}(l; k) - \underline{\Delta}(l; k)$. For example, in [Figure 2\(b\)](#), oscillation between the two hollow circles is uninformative under parameter value k , yet for $k' < k$, an oscillation of this size cannot be uninformative at the same public likelihood ratio because the difference in the net congestion thresholds is larger. As k vanishes, the difference between the threshold curves grows unboundedly at any interior public likelihood ratio ([Lemma 2](#)); hence, for small k , uninformative oscillation cannot occur perpetually except at very extreme public likelihood ratios given that incremental congestion effects are bounded ([Assumption 2](#)). This is depicted in [Figure 2\(b\)](#) by an uninformative oscillation of the same size occurring at a more extreme public likelihood ratio under k' than under k . [Theorem 1](#) below formalizes the idea into a result on learning with high probability at the no-congestion limit.

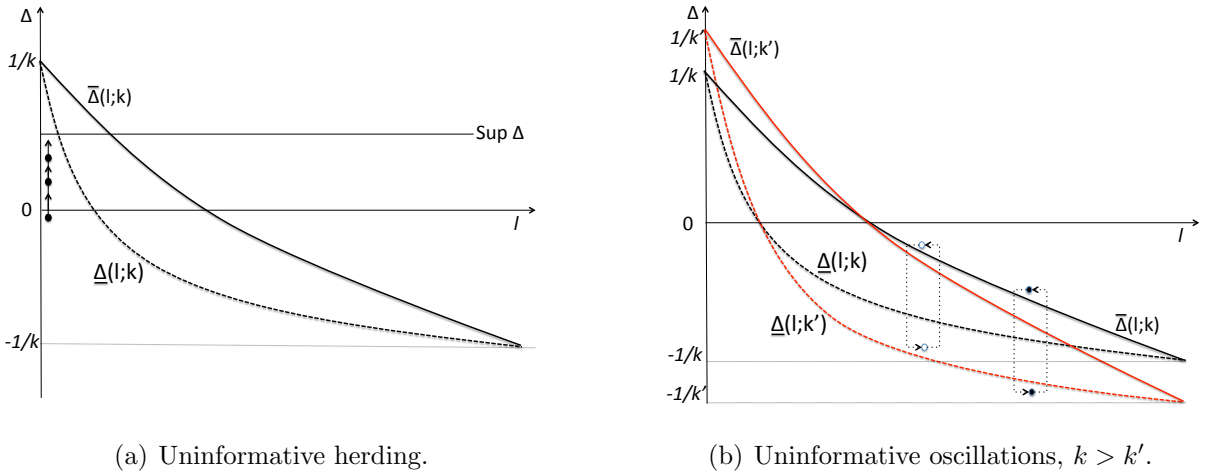


Figure 2 – Net-congestion costs and learning.

To substantiate the above discussion, consider the linear versions of [Example 1](#) and [Example 2](#). In the linear proportional-cost model, the incremental effect that any agent has on net congestion eventually vanishes. Hence, for any $k > 0$, perpetual uninformative oscillation cannot occur; the only reason that complete learning would not obtain is herding. When $k < 1$,

congestion costs can never become large enough to break a herd at sufficiently extreme public likelihood ratios. When $k \geq 1$, congestion costs eventually do become high enough to prevent herds at any (interior) public likelihood ratio. In the linear absolute-cost model, congestion costs can always get high enough to prevent a herd at any public likelihood ratio. In fact, since each agent's action has a constant incremental effect (in absolute value) on net congestion, agents will eventually oscillate uninformatively between the two actions on any sample path. Yet, as k vanishes, these uninformative oscillations can only occur at extreme public likelihood ratios.

We summarize the properties of asymptotic learning in the leading examples as follows.

Proposition 1. *Consider the leading examples.*

1. *In the linear proportional-costs model, there is (i) complete learning if $k \geq 1$; (ii) bounded learning if $k < 1$; and (iii) bounded learning in the no-congestion limit.*
2. *In the linear absolute-costs model, there is (i) bounded learning for any $k > 0$; but (ii) learning with high probability in the no-congestion limit.*
3. *In a non-linear proportional-costs model with an $f(\cdot)$ function whose range is $(-\infty, \infty)$, there is complete learning for any $k > 0$, and hence in the no-congestion limit as well.*
4. *In a non-linear absolute-costs model with an $f(\cdot)$ function whose range is bounded, there is bounded learning for all sufficiently small k and bounded learning in the no-congestion limit.*

(See the Appendix for a proof of this result and of all others not in the text.)

4.3 The no-congestion limit

We begin the formal analysis by studying learning in the no-congestion limit. As we will see, when $k \rightarrow 0$ the only reason that learning can stall is herding, which turns on properties of total congestion.

Definition 1. Total congestion is *bounded* if

$$\text{for any } a^i: -\infty < \inf \Delta(a^i) \text{ and } \sup \Delta(a^i) < +\infty.$$

Total congestion is *unbounded* if

$$\text{for any } a^i: \lim \Delta((a^i, 1, 1, \dots)) = +\infty \text{ and } \lim \Delta((a^i, -1, -1, \dots)) = -\infty.$$

In the leading examples ([Example 1](#) and [Example 2](#)), whether congestion is bounded or unbounded depends in each case on the range of the function $f(\cdot)$. In particular, the linear proportional-costs model has bounded total congestion whereas the linear absolute-costs model has unbounded total congestion. We note that although unbounded total congestion

and bounded total congestion are exclusive, they are not exhaustive because each property is required to hold for all action sequences.

Theorem 1. *If total congestion is unbounded, then there is learning with high probability in the no-congestion limit. If total congestion is bounded, then there is bounded learning in the no-congestion limit.*

Proof. The second statement of the theorem will be proved later as a consequence of [Theorem 3](#) and [Proposition 3](#) (see [Corollary 1](#)); we prove here the first statement of the theorem.

Step 1: We first claim that for any $x > 0$, $x \notin \text{Supp}[l_\infty^k]$ once k is small enough. Suppose not, per contra, for some $x > 0$. Then, for any $\varepsilon > 0$, there is a sequence of $k \rightarrow 0$ such that for each k , $\Pr[l_\infty^k \in B_\varepsilon(x)] > 0$, where $B_\varepsilon(x) := \{y : |x - y| < \varepsilon\}$. For $\varepsilon > 0$ small enough, there is an interval $[b_*(x, \varepsilon), b^*(x, \varepsilon)]$ with $\bar{b} > b^*(x, \varepsilon) > b_*(x, \varepsilon) > \underline{b}$ such that if a sample path has l_i^k and l_{i+1}^k both in $B_\varepsilon(x)$, then i 's private belief threshold—i.e., the lowest private belief that induces i to choose $a = 1$ (recall the monotonicity assured by [Lemma 1](#))—cannot fall within $[b_*(x, \varepsilon), b^*(x, \varepsilon)]$.¹⁵ Just as in the proof of [Lemma 2](#), $b^*(x, \varepsilon)$ can be mapped into a net-congestion threshold, $\Delta^*(x, \varepsilon; k) < \bar{\Delta}(x; k)$, through

$$1 + k\Delta^*(x, \varepsilon; k) = 2p(b^*(x, \varepsilon), x), \quad (3)$$

and analogously $b_*(x, \varepsilon)$ maps into a threshold $\Delta_*(x, \varepsilon; k) > \underline{\Delta}(x; k)$ through

$$1 + k\Delta_*(x, \varepsilon; k) = 2p(b_*(x, \varepsilon), x). \quad (4)$$

For l_i^k and l_{i+1}^k to both lie in $B_\varepsilon(x)$, it must be that $\Delta(a^{i-1}) \notin [\Delta_*(x, \varepsilon; k), \Delta^*(x, \varepsilon; k)]$; otherwise i 's private belief threshold would fall within $[b_*(x, \varepsilon), b^*(x, \varepsilon)]$. Combining (3) and (4) yields

$$k(\Delta^*(x, \varepsilon; k) - \Delta_*(x, \varepsilon; k)) = 2(p(b^*(x, \varepsilon), x) - p(b_*(x, \varepsilon), x)),$$

and hence

$$\text{as } k \rightarrow 0, \Delta^*(x, \varepsilon; k) - \Delta_*(x, \varepsilon; k) \rightarrow \infty. \quad (5)$$

Now, for any small enough $\varepsilon > 0$, let t be a time on some sample path such that $l_i^k \in B_\varepsilon(x)$ for all $i \geq t$; this is well-defined because $\Pr[l_\infty^k \in B_\varepsilon(x)] > 0$. Since $\Delta(a^i) \notin [\Delta_*(x, \varepsilon; k), \Delta^*(x, \varepsilon; k)]$ for all $i \geq t$, the fact that total congestion is unbounded implies that there must be an infinite set of agents, $I \subseteq \{t, t+1, \dots\}$ such that for any $i \in I$, $\Delta(a^i) \leq \Delta_*(x, \varepsilon; k)$ whereas $\Delta((a^i, 1)) \geq \Delta^*(x, \varepsilon; k)$. However, because of (5), bounded incremental congestion ([Assumption 2](#)) implies that once k is small enough, if $\Delta(a^i) \leq \Delta_*(x, \varepsilon; k)$ then $\Delta((a^i, 1)) < \Delta^*(x, \varepsilon; k)$, a contradiction.

¹⁵To confirm this, let \hat{b}_i denote i 's private belief threshold. If $a_i = 1$, then l_{i+1}^k is derived through Bayesian-updating using l_i^k and the event $b_i > \hat{b}_i$, hence for any $\hat{b}_i > \underline{b}$, there is a $\delta > 0$ such that $l_{i+1}^k \notin B_\delta(x)$. Since $a_i = -1$ works *mutatis mutandis*, it follows that given any small enough $\varepsilon > 0$, $l_{i+1}^k \in B_\varepsilon(x)$ only if \hat{b}_i is sufficiently close to either \underline{b} (if $a_i = 1$) or \bar{b} (if $a_i = -1$).

Step 2: We next claim that for any $\varepsilon > 0$ and $\delta > 0$, $\Pr[l_\infty^k > \varepsilon] < \delta$ for all small enough k . To prove this, fix any $\varepsilon > 0$ and $\delta > 0$. Let L be any number strictly larger than $1/\delta$. By Fatou's Lemma, $\mathbb{E}[l_\infty^k] \leq l_1^k = 1$, where the equality is from the neutral prior. This implies that $\Pr[l_\infty^k > L] < \delta$: if not, we would have $\mathbb{E}[l_\infty^k] \geq \delta L > 1$, a contradiction. The claim now follows from Step 1's implication that $\text{Supp}[l_\infty^k] \subseteq [0, \varepsilon] \cup [L, \infty)$ once k is small enough.

For any $\varepsilon > 0$, by applying the above claim to a sequence of $\delta \rightarrow 0$, it holds that $\lim_{k \rightarrow 0} \Pr[l_\infty^k \leq \varepsilon] = 1$; hence, there is learning with high probability in the no-congestion limit. Q.E.D.

Theorem 1 shows that whether learning is likely to occur when k gets small turns on whether total congestion is unbounded. The theorem can be interpreted as identifying a sense in which the bounded learning conclusion of the standard model without congestion (recall **Remark 1**) is fragile, but one should be clear in what sense our model does and does not converge to the standard model as $k \rightarrow 0$. To this end, let $\tilde{u}_i(a_i, \theta) = \mathbb{1}_{\{a_i = \theta\}}$ represent preferences in the model without congestion and $\tilde{u}_i^k(a^i, \theta) = u_i(a^i, \theta) = \mathbb{1}_{\{a_i = \theta\}} - kc(a^i)$ represent preferences under congestion factor $k > 0$. Then, for any i and $\varepsilon > 0$, there exists $\delta(i, \varepsilon) > 0$ such that if $k < \delta(i, \varepsilon)$ then for all θ and all a^i : $|\tilde{u}_i^k(a^i, \theta) - \tilde{u}_i(a_i, \theta)| < \varepsilon$. In other words, as $k \rightarrow 0$, our model converges pointwise across agents to the model without congestion. The convergence is not uniform, however, when total congestion is unbounded: the values of $\delta(i, \varepsilon)$ cannot be chosen independently of i . By contrast, there is uniform convergence when total congestion is bounded. Nevertheless, the case of unbounded total congestion seems pertinent for many economic applications in which immense congestion causes immense distress.

4.4 Complete learning for any $k > 0$

We next turn to studying properties of congestion costs that assure complete learning for an arbitrary $k > 0$. As suggested by the discussion in **Subsection 4.2**, one key condition is that total congestion must get significant enough in magnitude to prevent herds.

Definition 2. For any $k > 0$, total congestion *can get large* or *is large* if

$$\text{for any } a^i : \lim \Delta((a^i, 1, 1, \dots)) \geq 1/k \text{ and } \lim \Delta((a^i, -1, -1, \dots)) \leq -1/k.$$

Plainly, if total congestion can get large for some $k > 0$, then it also can get large for $k' > k$. In particular, total congestion is unbounded if and only if it can get large for all $k > 0$.

In addition to precluding herds, complete learning also requires that agents must not perpetually oscillate between actions without conveying information about their private beliefs.

Definition 3. Congestion has *no gaps* provided that for any $\varepsilon > 0$ and any non-convergent infinite action sequence (a_1, \dots) : if $S \subseteq \mathbb{R}$ is a bounded interval and I_S is an infinite set of agents such that $i \in I_S \iff \Delta(a^i) \in S$, then there is some i_S^* such that for any $i, j > i_S^*$ with $i, j \in I_S$ and for any $x \in (\Delta(a^i), \Delta(a^j))$, there exists $n > \max\{i, j\}$ such that $\Delta(a^n) \in (x - \varepsilon, x + \varepsilon)$.

Note that the no gaps condition is independent of k . While the condition may appear complicated, it has a fairly straightforward interpretation. To see this, assume that total congestion is bounded and pick any non-convergent infinite sequence of actions. Roughly, [Definition 3](#) requires that if we choose any two agents, i and j , far enough down the sequence, then the interval of net congestion costs $(\Delta(a^i), \Delta(a^j))$ can be arbitrarily finely “covered” by subsequent net congestion cost levels, in the sense that $\bigcup_{n > \max\{i, j\}} \Delta(a^n)$ creates an arbitrarily fine grid in that interval.¹⁶ The no-gaps condition is sufficient to ensure that agents cannot oscillate forever between the two actions without conveying any information.

Theorem 2. *If congestion has no gaps, then there is complete learning at any $k > 0$ for which total congestion can get large. Therefore, if congestion has no gaps and total congestion is unbounded, there is complete learning at the no-congestion limit.*

Proof. Fix some $k > 0$ and assume that total congestion can get large and congestion has no gaps. We will prove that $\text{Supp}[l_\infty^k] = \{0\}$. Suppose, per contra, that $x > 0$ and $x \in \text{Supp}[l_\infty^k]$. Following the logic developed in the proof of [Theorem 1](#) and using the notation introduced there for $\Delta^*(\cdot)$ and $\Delta_*(\cdot)$, we conclude that for any $\varepsilon > 0$ small enough, there must be a sample path of actions (a_1, \dots) and some time t such that

$$\text{for all } i \geq t: \text{ either } \Delta(a^i) \leq \Delta_*(x, \varepsilon; k) \text{ or } \Delta(a^i) \geq \Delta^*(x, \varepsilon; k). \quad (6)$$

Since large total congestion implies

$$\begin{aligned} \lim \Delta((a^i, -1, -1, \dots)) &\leq -1/k \leq \underline{\Delta}(x; k) \\ &< \Delta_*(x, \varepsilon; k) \\ &< \Delta^*(x, \varepsilon; k) < \overline{\Delta}(x, \varepsilon; k) \leq 1/k \leq \lim \Delta((a^i, 1, 1, \dots)), \end{aligned}$$

it follows that

$$|\{i : i \geq t \text{ and } \Delta(a^i) \leq \Delta_*(x, \varepsilon; k)\}| = |\{i : i \geq t \text{ and } \Delta(a^i) \geq \Delta^*(x, \varepsilon; k)\}| = \infty.$$

Thus, given any i^* , we can find $i, j > \max\{i^*, t\}$ such that $\Delta(a^i) \leq \Delta_*(x, \varepsilon; k)$ and $\Delta(a^j) \geq \Delta^*(x, \varepsilon; k)$. Furthermore, because of bounded increments ([Assumption 2](#)), (6) implies that there is some bounded interval, $S(x, \varepsilon; k) \supseteq [\Delta_*(x, \varepsilon; k), \Delta^*(x, \varepsilon; k)]$, such that $\Delta(a^i) \in S(x, \varepsilon; k)$ for all i . It then follows from the no-gaps property that for some $n > t$, $\Delta(a^n) \in (\Delta_*(x, \varepsilon; k), \Delta^*(x, \varepsilon; k))$; but this contradicts (6). *Q.E.D.*

In terms of the no-congestion limit, [Theorem 2](#) strengthens the positive conclusion of [Theorem 1](#) from learning with high probability to complete learning, but requires congestion to have no gaps. More importantly, [Theorem 2](#) can be applied to arbitrary $k > 0$.

¹⁶Since we are concerned with what happens far enough in the action sequence, it would be more accurate to call the property “*eventually* no gaps”, but we omit the “eventually” qualifier for brevity.

While congestion having no gaps is the essence of what drives complete learning, [Definition 3](#) is not always easy to verify directly in applications. Accordingly, we introduce the following stronger condition:

Definition 4. There is *vanishing incremental congestion* if for any $\varepsilon > 0$, any infinite sequence of actions (a_1, \dots) , and any bounded interval $S \subseteq \mathbb{R}$, there exists an i' such that if $i > i'$ and $\Delta(a^i) \in S$, then $|\Delta(a^{i+1}) - \Delta(a^i)| < \varepsilon$.

Intuitively, vanishing incremental congestion requires that in any infinite action sequence, eventually any two successive agents, i and $i + 1$, face net congestion costs that are arbitrarily similar. For this to hold, not only should agent i 's own action not affect net congestion by much, but moreover, the manner by which the actions of agents $1, \dots, i - 1$ affect i 's net congestion cost must be similar to how they affect agent $i + 1$'s. Thus, although reasonable in many contexts, vanishing incremental congestion rules out certain kinds of time-varying congestion costs that may nevertheless satisfy no gaps.

Proposition 2. *Vanishing incremental congestion implies that congestion has no gaps.*

The following observations show how [Proposition 2](#) can be used to apply [Theorem 2](#) to our leading examples. The [proportional-cost model](#) with a continuous $f(\cdot)$ satisfies vanishing incremental congestion: intuitively, eventually any one agent's action has a negligible effect on the fraction of agents who have chosen either action.¹⁷ Moreover, total congestion is large in this model if and only if $\sup f(\cdot) \geq 1/k$. In particular, total congestion gets large in the linear proportional-cost model if and only if $k \geq 1$. Similarly, total congestion is large in the [absolute-cost model](#) if and only if $f(\cdot)$ is unbounded above; hence, the linear version has large total congestion costs for all $k > 0$. However, the no gaps condition fails in the linear absolute-cost model because $\Delta((a^i, 1)) - \Delta(a^i) = \Delta(a^i) - \Delta((a^i, -1)) = 1$ for all a^i . On the other hand, it is straightforward to show that when $f(\cdot)$ is bounded, the absolute-cost model satisfies vanishing incremental congestion.

4.5 Bounded learning for any $k > 0$

Our final set of results on asymptotic learning derive sufficient conditions for bounded learning for arbitrary $k > 0$.

Definition 5. For any $k > 0$, congestion has *gaps* if there exists $C(k) > 0$ such that for any infinite action sequence (a_1, \dots) , there is some i^* such that for all $i > i^*$, $\Delta(a^i) \notin (1/k - C(k), 1/k) \cup (-1/k, -1/k + C(k))$.

¹⁷If the continuous function $f(\cdot)$ is unbounded, negligible changes in $\rho(\cdot)$, the fraction of agents who have chosen action 1, can still have a non-negligible effect on net congestion as $\rho(\cdot)$ approaches 0 or 1. However, [Definition 4](#) only requires incremental congestion to vanish along a sequence in which net congestion stays bounded. Formally, given any interval $[y_1, y_2] \subseteq [0, 1]$ (with $y_1 > 0$ if $\lim_{y \rightarrow 0} f(y) = -\infty$ and $y_2 < 1$ if $\lim_{y \rightarrow 1} f(y) = \infty$), any infinite sequence of actions (a_1, \dots) along which $\rho(a^i)$ stays within $[y_1, y_2]$ has the following property: given any $\varepsilon > 0$, there is i' such that for any a^i with $i > i'$, $\max\{f(\rho((a^i, 1))) - f(\rho(a^i)), f(\rho(a^i)) - f(\rho((a^i, -1)))\} < \varepsilon$. This follows from the continuity of $f(\cdot)$ and that eventually any one agent's choice has negligible effect on $\rho(\cdot)$.




Figure 3 – A sample path along which there is learning.

In words, the gaps condition precludes any infinite sequence of net congestion cost from converging to $1/k$ from below or to $-1/k$ from above. To see the intuition for why this implies bounded learning, consider [Figure 3](#), which depicts a sample path along which there is learning. As the figure suggests, any sample path with learning requires that there be a sequence of agents for whom net congestion cost converges to $1/k$ from below, contrary to the gaps condition. We will see momentarily how the gaps condition subsumes the intuitions provided in [Subsection 4.2](#) about bounded learning when either congestion costs on one action never get large enough ([Proposition 3](#) below) or when the incremental effect of an agent’s action on successors’ net congestion costs is never negligible ([Theorem 3](#) below).

Remark 2. Except in degenerate cases—such as when there are no congestion effects—the gaps condition and the no gaps condition are incompatible. In particular, if total congestion can get large, then both conditions cannot hold simultaneously. ||

A simple sufficient condition for congestion to have gaps for all $k > 0$ is that

$$\text{the range of } \Delta(\cdot) \text{ has no finite limit point.} \tag{7}$$

Clearly, the [linear absolute-cost model](#) satisfies (7) because in this case $\Delta(\cdot) \in \mathbb{Z}$.

Theorem 3. *For any $k > 0$, if congestion has gaps then there is bounded learning. Furthermore, if for all $k > 0$ small enough it holds that congestion has gaps and the constant $C(k)$ in [Definition 5](#) can be chosen such that $1/k - C(k)$ is bounded above, then there is bounded learning in the no-congestion limit.*

Proof. For the first statement, fix any $k > 0$ and assume that congestion has gaps. Let \hat{l} be

defined by $\underline{\Delta}(\hat{l}; k) = 1/k - C(k)$, where $C(k) > 0$ is from [Definition 5](#); without loss, we may take $1/k - C(k) > 0$. Since $\underline{\Delta}(\cdot; k)$ is strictly decreasing from $1/k$ to $-1/k$ ([Lemma 2](#)), \hat{l} is well defined. Pick an arbitrary sample path of actions (a_1, \dots) . Since $1/k > \bar{\Delta}(l; k) > \underline{\Delta}(l; k)$ for all $l > 0$, it follows that if $l_i^k < \hat{l}$ then i plays uninformatively; hence there cannot be an i such that $l_i^k < \frac{1-p(\bar{b}, \hat{l})}{p(\bar{b}, \hat{l})} \hat{l}$, where the latter term is strictly larger than 0. Since the sample path was arbitrary, it follows that there is bounded learning.

For the second statement, assume there is some $z \in (0, \infty)$ such that $1/k - C(k) < z$ for all small enough $k > 0$. Define \hat{l} by $\underline{\Delta}(\hat{l}; k) = z$. The same argument as above can be used to conclude that for any $k > 0$ small enough, $\min \text{Supp}[l_\infty^k] \geq \frac{1-p(\bar{b}, \hat{l})}{p(\bar{b}, \hat{l})} \hat{l}$, which implies bounded learning in the no-congestion limit. Q.E.D.

Remark 3. For a given $k > 0$, the gaps condition yields bounded learning because [Definition 5](#) requires the same constant, $C(k)$, to apply to all infinite action sequences. If, instead, the constant could depend on the action sequence, then we would only have the weaker conclusion that beliefs almost surely do not converge to the truth, i.e. $\Pr(l_\infty^k = 0) = 0$. ||

[Theorem 3](#) can be used to deduce what happens when total congestion cost is small in the following sense:

Definition 6. For any $k > 0$, total congestion is *small* if $-1/k < \inf \Delta(\cdot)$ and $\sup \Delta(\cdot) < 1/k$.

Proposition 3. *For any $k > 0$, if total congestion is small then congestion has gaps. If total congestion is bounded then congestion has gaps for all $k > 0$ small enough; furthermore, the constant $C(k)$ in [Definition 5](#) can be chosen such that $1/k - C(k)$ is bounded above.*

Combining [Theorem 3](#) and [Proposition 3](#) yields:

Corollary 1. *For any $k > 0$, if total congestion is small then there is bounded learning. If total congestion is bounded then there is bounded learning in the no-congestion limit.*

Recall from [Remark 1](#) that in the benchmark model with no congestion, although herds occur almost surely, information cascades need not. In other words, bounded learning does not imply information cascades. Moreover, at any point where a cascade arises, so too does a herd. By contrast, with congestion costs, information cascades need not usher in herds. The next result describes properties of the net congestion cost that guarantee the onset of an information cascade and simultaneously rule out herding.

Proposition 4. *If total congestion is unbounded and [\(7\)](#) holds, then for any $k > 0$ there is almost surely an information cascade without a herd.*

Suppose [\(7\)](#) holds and total congestion is unbounded—as is true, for example, in the linear absolute-cost model. Then, [Theorem 3](#) and [Proposition 4](#) imply that for any $k > 0$, learning is bounded and an information cascade will almost surely arise. It is important to recognize, however, that as $k \rightarrow 0$, the set of public beliefs at which learning stalls and cascades arise changes. Since there is learning with high probability in the no-congestion limit ([Theorem 1](#)), the public belief can only settle when it is arbitrarily confident as k vanishes.

4.6 The strength of public beliefs with large total congestion

In the [introduction](#), we mentioned how in the [linear-absolute cost](#) model even small incremental congestion costs can substantially affect asymptotic public beliefs in comparison to the standard no-congestion model. We now further illustrate how congestion improves learning for a broader class of congestion costs as well as different distributions of private beliefs. Say that $b \geq 1/2$ is an upper (resp. lower) bound on the *strength* of the asymptotic public belief if the asymptotic public belief must lie in (resp. outside) $[1 - b, b]$. It is straightforward that when $k = 0$, an upper bound on the strength of the public belief is $p\left(\bar{b}, \frac{1-\bar{b}}{\bar{b}}\right)$.

Let $C > 0$ denote a bound on eventual incremental congestion cost:¹⁸

$$\text{for any } (a_1, \dots): \text{ there exists an } i' \text{ such that if } i > i', \text{ then } |\Delta(a^{i+1}) - \Delta(a^i)| \leq C, \quad (8)$$

and assume that total congestion can get large. In order for learning to stall at a public likelihood ratio l , an infinite subsequence of players must switch uninformatively between $a = 1$ and $a = -1$; (8) implies that this can only happen if $\bar{\Delta}(l; k) - \underline{\Delta}(l; k) \leq C$. As discussed in [Lemma 2](#) and its proof, $\bar{\Delta}(l; k) - \underline{\Delta}(l; k)$ is simply a multiple of $1/k$ and $\bar{\Delta}(l; k) - \underline{\Delta}(l; k)$ is quasi-concave in l . Consequently, so long as $k \times C$ is sufficiently small, there will be two public likelihood ratios at which

$$\bar{\Delta}(l; k) - \underline{\Delta}(l; k) = C, \quad (9)$$

and the solutions only depend on $k \times C$. Furthermore, for any l that lies in between the two solutions, $\bar{\Delta}(l; k) - \underline{\Delta}(l; k) > C$. It can be verified that if $\underline{b} = 1 - \bar{b}$, the two solutions to [Equation 9](#) are symmetric in the sense that the corresponding public beliefs add up to one. Consequently, when $\underline{b} = 1 - \bar{b}$ and $k \times C$ is sufficiently small, a *lower* bound on the strength of the asymptotic public belief is the public belief corresponding to the smaller solution of [Equation 9](#); the distribution of private beliefs does not matter beyond the value of $\underline{b} = 1 - \bar{b}$.

[Table 1](#) illustrates these lower bounds for some parameter values, and the corresponding *upper* bounds on the strength of public beliefs in the no-congestion model. We see that the smaller are incremental congestion costs in the sense of $k \times C$, the sharper is the lower bound on the strength of the asymptotic public belief (so long as total congestion can still get large, e.g. it is unbounded). The discrepancy between the no-congestion and small-incremental-congestion models is starker when private beliefs are weaker, which are cases generally viewed as more relevant for observational learning. Consider $\bar{b} = \frac{5}{9}$: the strength of the asymptotic public belief in the no-congestion model is bounded from above by approximately 0.61. By contrast, with $C = 1$, tiny incremental congestion costs of $k = 0.001$ —an agent in the linear absolute-cost model, for instance, is deterred from choosing an action known to be good only once its queue is 1,000-people longer—ensure that limiting public beliefs must exceed approximately

¹⁸By [Assumption 2](#), incremental congestion is bounded. What matters for the current exercise is only a (possibly tighter) bound on eventual incremental congestion. In the linear absolute-cost model, $C = 1$.

	$\bar{b} = \frac{2}{3}$	$\bar{b} = \frac{3}{5}$	$\bar{b} = \frac{4}{7}$	$\bar{b} = \frac{5}{9}$
$k = 0$	0.8	0.6923	0.64	0.6098
$k \times C = 0.1$	0.9649	0.9352	0.9046	0.8718
$k \times C = 0.01$	0.9967	0.9940	0.9913	0.9888
$k \times C = 0.001$	0.9997	0.9994	0.9991	0.9989

Table 1 – Lower bounds on strength of asymptotic public belief when $\underline{b} = 1 - \bar{b}$ and eventual incremental congestion is bounded by C . Upper bounds for $k = 0$.

0.999 (or lie below 0.001) if total congestion can get large. In this sense, even in cases where learning is incomplete for small k , congestion costs may substantially improve the accuracy of asymptotic public beliefs. It bears repeating, however, that this analysis ignores the speed of convergence, which may be slower with lower k .

5 Applications

5.1 Congestion through transfers

Having addressed the question of asymptotic learning, it is natural to next ask what fraction of agents asymptotically choose the superior action, i.e. the action that matches the state. Aside from being intrinsically interesting and potentially empirically observable, this statistic is generally relevant for any reasonable measure of welfare.¹⁹ Moreover, we will see that it has useful application.

Proposition 5. *Assume (i) total congestion can get large; (ii) there is vanishing incremental congestion; and (iii) $\Delta(a^i) = g(\rho(a^i))$ for some continuous and injective function $g : [0, 1] \rightarrow \mathbb{R}$. Then, ex-ante, the proportion of agents who asymptotically choose the superior action almost surely is $1/2 + (1/2)[g^{-1}(1/k) - g^{-1}(-1/k)]$.*

Recall that the first two hypotheses of [Proposition 5](#) imply that there is complete learning for any $k > 0$. This does not guarantee that ρ , the proportion of agents choosing action 1 or the action frequency, converges almost surely. Augmenting the proposition’s third hypothesis does ensure this, as established in the proof.²⁰ It is then straightforward to see that, almost

¹⁹This statistic can also be compared with its counterpart in the standard model without congestion. Comparing equilibrium utilities (even asymptotically) across the two settings is less compelling. If one views agents’ utility functions as just representing their preferences, then the comparison is not very meaningful since preferences differ. Furthermore, congestion costs affect an agent’s behavior entirely through the difference in congestion cost an agent faces between the two actions. This means that behavior in our model is isomorphic to behavior in a model with congestion “benefits” where the history-dependent benefit of taking action a is the cost we subtract from taking action $-a$. Framing congestion as a cost or as a benefit will clearly affect any welfare conclusion drawn by comparing utilities across models.

²⁰Even though the gaps condition would suffice for complete learning (given large total congestion), the stronger

surely, $\rho \rightarrow g^{-1}(1/k)$ when $\theta = 1$ and $1 - \rho \rightarrow 1 - g^{-1}(-1/k)$ when $\theta = -1$, which explains the formula in the proposition.

Proposition 5 implies that the proportion of agents who choose the superior action converges to one in some environments. Consider, for instance, the [linear proportional-cost model](#) when total congestion is large, i.e. $k \geq 1$. Applying the expressions mentioned above, the proportion of agents choosing the superior action in this example is $(1 + k)/2k$ regardless of the state. Thus, when $k = 1$, the fraction of agents who eventually choose the superior action is one. The higher is k , the lower is the proportion of agents who eventually choose the superior action, and this proportion converges to $1/2$ as $k \rightarrow \infty$.²¹

An interesting corollary is that in the canonical observational learning model without congestion, a social planner could use a simple transfer scheme to ensure that, asymptotically, all agents will choose the superior action. This can be achieved by using transfers to induce congestion costs that take the form of the linear proportional-cost model with $k = 1$. Specifically, a social planner could require agent i to pay an amount $\rho(a_i^{i-1})$ to take action $a_i = 1$ and pay $1 - \rho(a_i^{i-1})$ to take action $a_i = -1$. Agent i 's transfer could be redistributed arbitrarily as a form of subsidy to subsequent agents independent of their choices.²² By [Proposition 5](#), this simple transfer scheme causes agents to internalize their information externality in such a way that leads to asymptotic efficiency.

5.2 Congestion through market prices

A simple application of our results is to a market setting where congestion cost is induced by a price mechanism. There are two products, A and B . It is known that one product is of high quality and the other of low quality, but consumers do not know which is which. We represent A (resp. B) being the high-quality product by the state $\theta = 1$ (resp. $\theta = -1$). Gross of price, the value of the high-quality product to any consumer is 1 while the low-quality product is 0. Denote the decision to purchase product A by $a_i = 1$ and the decision to purchase product B by $a_i = -1$. Consumers make purchase decisions sequentially and each consumer observes the history of purchase decisions.

Assume that product pricing is as follows: after any history of consumer purchases, the price of each product equals a constant $k \leq 1$ times its expected value conditional on the public information. When $k = 1$, this model is similar to the leading example of [Avery and](#)

requirement of vanishing incremental congestion is used in the argument establishing almost-sure convergence of the action frequency. Note that action-frequency convergence is not an issue in the standard model without congestion, because there a herd ensues almost surely.

²¹The [linear absolute-cost model](#) is not covered by [Proposition 5](#), because it does not lead to complete learning. However, it is easy to reason directly in this case that the action frequency converges almost surely to $1/2$, because agents eventually oscillate in succession between the two actions. Interested readers may refer to earlier versions of this paper that contain a more detailed analysis of action-frequency convergence when there is incomplete learning.

²²Redistribution to preceding agents is not straightforward because it would introduce strategic forward-looking considerations into all agents' choices.

Zemsky (1998), but with an arbitrary distribution of signals of bounded informativeness. The parameter range $k < 1$ can be interpreted as a reduced-form representation of competition between sellers who face zero marginal costs of production. At the extreme, the case of $k = 0$ corresponds to perfect or Bertrand competition.

Using our notation for public beliefs, the price of good A after history a^i is $kq(a^i)$, while the price of product B is $k(1 - q(a^i))$. Thus, ignoring indifference as usual, consumer $i + 1$ with private beliefs b chooses $a_{i+1} = 1$ (i.e. buy A) if and only if $p(b, q(a^i)) - kq(a^i) > 1 - p(b, q(a^i)) - k(1 - q(a^i))$, or equivalently,

$$p(b, q(a^i)) > \frac{1}{2} + \frac{k}{2} (2q(a^i) - 1). \quad (10)$$

Even though there is no explicit congestion cost, the (endogenous) price has a similar effect. Indeed, for arbitrary $k \leq 1$, we can define the net congestion cost

$$\Delta(a^i) := 2q(a^i) - 1 = \frac{1 - l(a^i)}{1 + l(a^i)}, \quad (11)$$

so that the posterior belief threshold of our general model (Lemma 1), $p^*(\cdot) = 1/2 + k\Delta(\cdot)/2$, coincides with threshold implied by (10).²³ Note that $\Delta(\cdot)$ defined by (11) satisfies Assumption 1 and Assumption 2. We now consider two cases:

First consider $k = 1$. Plainly, for any a^i , there is $\varepsilon > 0$ such that $p(\bar{b}, l(a^i)) > q(a^i) + \varepsilon$ and $p(\underline{b}, l(a^i)) < q(a^i) + \varepsilon$. Alternatively, according to (11), for any $l > 0$ there is $\varepsilon > 0$ such that if $l(a^i) = l$ then $\underline{\Delta}(l; 1) + \varepsilon < \Delta(a^i) < \bar{\Delta}(l; 1) - \varepsilon$. This implies that at any interior public likelihood ratio, the informative content of an agent's action is always bounded away from zero: for any $l > 0$ there is $\varepsilon > 0$ such that if $l(a^i) = l$ then for any a_{i+1} , $|l(a^{i+1}) - l(a^i)| > \varepsilon$. Consequently, there is complete learning: the asymptotic public likelihood ratio can never settle at an interior point.

Now consider $k < 1$. Since (11) implies $\Delta(a^i) \in (-1, 1)$ for any a^i , the setting satisfies small total congestion for all $k < 1$. Hence, Corollary 1 implies that there is bounded learning for all $k < 1$ (and bounded learning in the limit of $k \rightarrow 0$). It is not hard to verify that on any sample path, there will be a herd in finite time.

Although we do not pursue it formally here, we could modify Avery and Zemsky (1998) in a different direction by assuming that the price updates after every $N > 1$ trades instead of after every trade, perhaps because some technological constraint prevents instant price-updating. This does not alter the aforementioned properties of asymptotic learning for either $k = 1$ or $k < 1$. However, even when $k = 1$, although there is complete learning, it holds that for sufficiently large $N > 1$, there is a strictly positive probability that an infinite number of agents behave uninformatively; this contrasts with $N = 1$ wherein every agent behaves infor-

²³Even though the function $\Delta(\cdot)$ defined by (11) depends indirectly on the strategies of the agents (because they affect the public belief), one can proceed recursively from agent one onward and just substitute this out.

matively.²⁴ Intuitively, within every block of N agents after a price update, the environment is analogous to the standard model without congestion: the first agent will act informatively because of the price update preceding him, whereas with sufficiently many agents in the block, there is a positive probability of a “temporary information cascade” within that block.

Another natural variation is to limit prices to lie on a grid, e.g. pounds and pence. As this restriction satisfies [gaps](#), [Theorem 3](#) implies that there is bounded learning in this setting.

5.3 Congestion through queuing

As a final application, we modify [Example 2](#) to study a queuing problem with stochastic service. In the *queuing model with constant unobservable service rate* $\delta \in [0, 1)$, $c(a^1) = 0$ and for each $i \geq 2$,

$$c(a^i) = \sum_{j=1}^{i-1} \mathbb{1}_{\{a_j=a_i\}} \delta^{i-j-1}. \quad (12)$$

Thus, congestion depends more strongly upon the choices made by recent predecessors than those made by more distant predecessors. If $\delta = 0$, every agent cares only about his immediate predecessor’s action.²⁵ This cost function naturally arises when agents suffer disutility from waiting for service under the constant service rate, where in every period, each previously un-served agent is served with probability δ . When actual service to preceding agents is unobservable, each agent computes the expectation over the number of predecessors who chose his action but have not yet been served. For instance, agent i may have observed whether each agent $j < i$ entered restaurant $a = -1$ or restaurant $a = 1$ through the front door. However, he may not know whether j was served and exited through the restaurant’s unobservable back door or remains in the restaurant. If agents are risk neutral and believe that in every period t , any agent who has entered a restaurant prior to t exits with probability $1 - \delta$, then each i faces a congestion cost given by [\(12\)](#).

From [\(12\)](#), one computes that for any a^i , $\Delta((a^i, 1)) = 1 + \delta\Delta(a^i)$ and $\Delta((a^i, -1)) = \delta\Delta(a^i) - 1$. Furthermore, for any a^i , $\lim \Delta(a^i, 1, 1, \dots) = \frac{1}{1-\delta}$, and $\lim \Delta(a^i, -1, -1, \dots) = -\frac{1}{1-\delta}$. Hence, total congestion is bounded. Total congestion can get large when $k \geq 1 - \delta$, whereas it is small when $k < 1 - \delta$. It follows immediately from [Theorem 1](#) that:

Corollary 2. *In the queuing model with constant unobservable service rate $\delta \in [0, 1)$, there is bounded learning in the no-congestion limit.*

One can also deduce properties of asymptotic learning when $k \geq 1 - \delta$. For example, since $\Delta(\cdot) \in \{-1, 1\}$ when $\delta = 0$, congestion has gaps in this case; hence, [Theorem 3](#) implies that there is bounded learning no matter the value of $k > 0$ when $\delta = 0$.

²⁴In fact, given any $x < 1$, N can be chosen large enough to make the probability of an infinite number of agents behaving uninformatively be strictly larger than x .

²⁵If δ were equal to one, this would be the [linear absolute-cost model](#).

We next turn to a variation of the problem where service is observable. In the restaurant context, whereas above diners departed through an unobservable back door, here they depart through the observable front door. Define the Bernoulli random variable $S_j(t)$ to equal 1 with probability δ when $t = j + 1$, and for each $t \geq j + 1$,

$$\Pr[S_j(t + 1) = 1 \mid S_j(t) = 1] = \delta, \quad \Pr[S_j(t + 1) = 1 \mid S_j(t) = 0] = 0.$$

In words, $S_j(t) = 1$ indicates that agent j remains unserved through period t , while $S_j(t) = 0$ indicates that he has been served. Agent i faces the congestion cost

$$c(a^i) = \sum_{j=1}^{i-1} \mathbb{1}_{\{a_j=a_i\}} S_j(i).$$

In words, agent i pays 1 for every unserved predecessor who chose the same action as i . Note that the cost faced by any agent is now stochastic.

Observe that for any $\delta > 0$ and any $k > 0$, congestion costs get arbitrarily large with probability one (e.g., when $S_j(\cdot) = 1$ for sufficiently many periods in a row). Using this fact, one can show that:

Corollary 3. *In the queuing model with constant observable service rate $\delta \in (0, 1)$, there is learning with high probability in the no-congestion limit.*

We omit the proof because it follows the same logic as that of [Theorem 1](#), with straightforward modifications to account for the stochastic costs. The contrast between [Corollary 2](#) and [Corollary 3](#) reiterates the theme that what is important for learning is not that every agent must play informatively, but rather that it is always inevitable that some agent in the future will do so.

6 Outside Options

Up until this point, we have modeled agents as necessarily having to choose between actions with uncertain payoffs that each suffers from congestion costs. In some applications, agents may also have some outside option, which provides a known payoff that is not affected by congestion—“dine at home”, so to speak, or “balk both queues”. If the congestion cost that agent i faces on both action -1 and action 1 is sufficiently large—as it necessarily is eventually, for example, in the [linear absolute-cost model](#) for any $k > 0$ —then regardless of i ’s posterior belief, he will prefer to take the outside option. This raises a question of whether our positive results on asymptotic learning rely on the absence of an un-congested outside option (or the violation of individual rationality). In this section, we show that this is not the case, under reasonable conditions. The logic of the analysis will also suggest how our results can be extended to an arbitrary number of actions with state-dependent payoffs and congestion costs.

We augment the model of [Section 2](#) by adding an outside action, denoted $a = 0$, that delivers a fixed payoff of zero to all agents who choose it: $u_i((a^{i-1}, 0), \theta) = 0$ for all i , a^{i-1} , and θ . Setting the outside option's payoff to zero entails a substantive assumption that if there were no congestion costs on either action -1 or 1 , then the outside option would be dominated; this is consistent with ignoring outside options in the standard model without congestion. Recall that without an outside option, the net congestion cost, $\Delta(a^i)$, was all that mattered to agent i besides his posterior belief. With the outside option, the actual level of congestion cost on each action $a_i \in \{-1, 1\}$, $c(a^i)$, also matters. Indeed, $c(a^i)$ is the analog of $\Delta(a^i)$ vis-à-vis the choice between action $a_i \in \{-1, 1\}$ and the outside option.

We will maintain the following generalization of bounded increments ([Assumption 2](#)):

Assumption 3 (Strengthened bounded increments). $\sup_{a^i, a, a'} |c((a^i, a)) - c((a^i, a, a'))| < \infty$.

In words, this requires that there is a bound on how much any agent's action can affect his successor's congestion cost, no matter the history.²⁶

Recall that our result of learning with high probability in the no-congestion limit ([Theorem 1](#)) relied on unbounded total congestion ([Definition 1](#)). Such a result in the presence of the outside option requires, in addition, the analogous property to also hold for $c(\cdot)$. To state the property, let $\neg a \in \{-1, 0, 1\} \setminus \{a\}$ denote an action different from a .

Definition 7. There is *full decay* if for any a^i and any $a \in \{-1, 1\}$: $\lim_{N \rightarrow \infty} c((a^i, \underbrace{\neg a, \neg a, \dots, \neg a}_N, a)) = 0$.²⁷

Full decay captures the notion that, eventually, if no one ever takes an action $a \in \{-1, 1\}$, then congestion on action a vanishes. It implies that there can never be a herd on the outside option, for if all agents start choosing the outside option at any point, then eventually an agent would find a different action strictly preferable, no matter his belief about the state. When combined, full decay and unbounded total congestion imply that if all agents start taking some action $a \in \{-1, 0, 1\}$ then, eventually, action a becomes dominated by another action no matter an agent's belief.

Theorem 4. *In the model with the outside option, if total congestion is unbounded and there is full decay, then there is learning with high probability in the no-congestion limit.*

Although the analysis with the outside option is more intricate than in the baseline model, the intuition underlying [Theorem 4](#) is similar to that of [Theorem 1](#): for small k , given that incremental congestion is bounded, the conjunction of full decay and unbounded total congestion imply that learning can only stall at very extreme public likelihood ratios. [Theorem 4](#) points

²⁶It would be natural to also assume a generalization of monotonicity ([Assumption 1](#)), but we do not need it for the formal result we will prove in this section.

²⁷It bears emphasis that the actions within the subsequence that occurs “ N times” are not assumed to be the same; rather, the subsequence just cannot contain a .

out that to obtain learning with high probability in the no-congestion limit, it is sufficient that the congestion cost on any action, in the sense of $k \times c$, be able to grow above 1 (the normalization of the maximum gross payoff from taking action 1 or -1). This requirement certainly does not imply that expected payoffs become arbitrarily poor on the equilibrium path; rather, they are always non-negative in the presence of the outside option.

We view full decay as a reasonable economic requirement. It is satisfied by any version of the [proportional-cost model](#) in which $f(0) = 0$. While the [linear absolute-cost model](#) violates the condition—and indeed, it is intuitive in this case that herding on the outside option is an obstacle to asymptotic learning even as k vanishes—a number of natural variants of this specification would satisfy full decay. For example, motivated by the queuing application with constant observable service rate, suppose that consumers are served in first-in-first-out order, but it takes $T \geq 1$ periods to process each consumer. The congestion cost to a consumer from joining queue $a \in \{-1, 1\}$ is equal to the amount of time before he finishes being served. This setting satisfies both bounded increments (per [Assumption 3](#)) and full decay. If $T = 1$, then $c(a^i) = 1$ for all i and a^i , hence $\Delta(a^i) = 0$ for all a^i , and there is bounded total congestion. On the other hand, for any $T > 1$, there is unbounded total congestion, hence [Theorem 4](#) applies.

While we have focussed above only on the no-congestion limit, it is also possible to derive analogous results to those of [Theorem 2](#) and [Theorem 3](#) for arbitrary $k > 0$ under analogous assumptions.

7 Conclusion

This paper has studied the role of congestion costs in an otherwise standard model of observational learning. Congestion costs capture situations in which an agent’s payoff from choosing an option decreases when more of his predecessors choose that option. This feature arises naturally in markets either through changing prices or market-share effects, and in other environments where, for example, costs may stem from delays in service or reduced benefits to conformists.

Our analysis sheds light on how different forms of congestion costs impact long-run learning. By parameterizing the marginal rate of substitution between the benefits from choosing the superior action and incurring congestion costs by $k > 0$, we have provided results for both arbitrary k and the no-congestion limit when k vanishes. While one might conjecture that the lessons from the literature without congestion costs would carry over to a model where each agent has only an infinitesimal effect on congestion costs, our results emphasize that this depends crucially on whether *total* congestion costs remain bounded or not. In many applications, one would posit that no matter the value of k , agents would eventually be willing to choose an action that is known to be inferior if there is enough congestion on the other actions. In such cases, asymptotic learning essentially obtains as k vanishes, in contrast to the case of $k = 0$ (with bounded private beliefs).

While we have focussed on settings where an agent finds an action more attractive when it has heretofore been rare, our methods can also be used to analyze situations where the direction of the externality is reversed, such as when agents have a taste for conformity. Formally, this just requires reversing the inequalities in the monotonicity assumption ([Assumption 1](#)). It is not hard to check that in this case, given bounded private beliefs, for any $k > 0$ there is bounded learning and almost surely a herd.²⁸

We have focussed on the case of bounded private beliefs, but the analysis can also be extended to unbounded private beliefs. In that case, complete learning obtains under small total congestion, just as in the standard model without congestion ([Smith and Sørensen, 2000](#)).²⁹ On the other hand, with large total congestion, complete learning may fail. For example, in the [linear absolute-cost model](#) (which has large total congestion, no matter the value of $k > 0$), there is bounded learning when k is sufficiently large, even when private beliefs are unbounded. Nevertheless, there is still learning with high probability in the no-congestion limit, just as under bounded private beliefs.

A key assumption in this paper is that only past actions affect the payoff of an agent. There are, of course, situations where future actions also matter. Consider the case of betting, where each bettor i bets on which of two teams will win a game. In some fixed-odds systems used by bookmakers, each bettor receives odds that depend on how many prior bettors have bet on each team, consistent with the backward-looking congestion costs of our model. In parimutuel betting, however, each bettor receives odds that depend on how many agents have chosen each team by the close of the betting pool. In this system, a bettor must consider not only his beliefs about the superior action and his predecessors' choices, but potentially also how his action influences the bets of his successors. Extending our analysis to such environments is a challenging but promising area of further research.³⁰

²⁸Under mild conditions one can also show that there is almost surely an information cascade. For example, the following condition would be sufficient: there exists $\varepsilon > 0$ such that for any $(a^i, 1, 1, \dots)$ and any $j > i$, $\Delta(a^j) > \varepsilon$ (and analogously if there is a herd on action -1), where $\Delta(a^j)$ is now the net benefit that individual j derives from taking action 1 rather than -1 .

²⁹If private beliefs are unbounded we have that $\bar{\Delta}(l; k) = 1/k$ and $\underline{\Delta}(l; k) = -1/k$ for all l , and hence an agent's action is always informative under small total congestion.

³⁰[Koessler et al. \(2008\)](#) make some progress with characterizing equilibrium behavior for sequential parimutuel betting with a very small number of bettors.

A Appendix

Proof of Proposition 1. Part 1(i) follows from [Theorem 2](#); Part 1(ii) from [Corollary 1](#); Part 1(iii) from [Theorem 1](#); Part 2(i) from [Theorem 3](#); Part 2(ii) from [Theorem 1](#); Part 3 from [Theorem 2](#); and Part 4 from [Corollary 1](#). Q.E.D.

Proof of Proposition 2. Fix any non-convergent infinite sequence of actions (a_1, \dots) and any $\varepsilon > 0$. Pick any bounded interval $S \subseteq \mathbb{R}$ and any infinite set of agents $I_S = \{n_1, n_2, \dots\}$ such that $i \in I_S \iff \Delta(a^i) \in S$; if such an S and I_S do not exist then we are trivially done. From the definition of vanishing incremental congestion, there are two exclusive and exhaustive possibilities: either (i) $x^* = \lim_{k \rightarrow \infty} \Delta(a^{n_k})$ exists, or (ii) there is a non-singleton closed interval $S' \subseteq S$ such that for any $\tilde{S} \subsetneq S'$ and any j , there is some $i > j$ such that $\Delta(a^i) \in S' \setminus \tilde{S}$, i.e. S' is the minimal set such that among agents in I_S , $\Delta(\cdot)$ is eventually in S' .

Assume case (i). This implies that there is some i^* such that $\Delta(a^i) \in (x^* - \varepsilon/2, x^* + \varepsilon/2)$ for all $i > i^*$ with $i \in I_S$. This implies that for any $i, j > i^*$ with $i, j \in I_S$, and any $x \in (\Delta(a^i), \Delta(a^j))$, it holds that for all $n > \max\{i, j\}$ with $n \in I_S$, $\Delta(a^n) \in (x - \varepsilon, x + \varepsilon)$. This satisfies the requirement of no gaps.

Now consider case (ii). Let i^* be any time such that $\Delta(a^i) \in S'$ for all $i > i^*$ and $i \in I_S$. By definition of S' , vanishing incremental congestion implies that for any $x \in S'$, there must be an infinite set of agents, $I' \subseteq I_S$, such that for all $i \in I'$, $\Delta(a^i) \in (x - \varepsilon, x + \varepsilon)$. This implies that the requirement of no gaps is satisfied because for any $i, j > i^*$ with $i, j \in I_S$, $\Delta(a^i) \in S'$ and $\Delta(a^j) \in S'$. Q.E.D.

Proof of Proposition 3. For the first statement, observe that for any $k > 0$, $C(k) = 1/k - \sup \Delta(a^i)$ verifies [Definition 5](#) when $\sup \Delta(a^i) < 1/k$. The second statement follows because under bounded total congestion, the same construction works for all $k > 0$ small enough, and in this case $1/k - C(k)$ is bounded above. Q.E.D.

Proof of Proposition 4. Fix $k > 0$ and pick any sample path (a_1, \dots) ; let z be public likelihood ratio limit on this sample path, which exists a.s. As total congestion is unbounded, the sample path includes an infinite number of each action. Assume, to contradiction, that there is no cascade. Then there is an infinite set of agents, I' , who all take the same action and whose actions are informative. Without loss, assume that $i \in I' \implies a_i = 1$; the argument proceeds *mutatis mutandis* in the other case. By the same logic used in proving [Theorem 1](#), it follows that for all small enough $\varepsilon > 0$, there exists i_ε such that for all $i > i_\varepsilon$, $\Delta(a^i) \notin (\underline{\Delta}(z; k) + \varepsilon, \overline{\Delta}(z; k) - \varepsilon)$. By the continuity of $\underline{\Delta}(\cdot; k)$ and $\overline{\Delta}(\cdot; k)$, it further follows that for all small enough $\varepsilon > 0$,

$$\text{there is } i_\varepsilon \text{ such that } [i > i_\varepsilon \text{ and } i \in I'] \implies \Delta(a^{i-1}) \in (\underline{\Delta}(z; k) - \varepsilon, \underline{\Delta}(z; k) + \varepsilon).$$

In view of (7), the above condition can hold for $\varepsilon > 0$ small enough only if

$$\text{there is } i' \text{ such that } [i > i' \text{ and } i \in I'] \implies \Delta(a^{i-1}) = \underline{\Delta}(z; k).$$

But this implies that eventually agents in I' behave uninformatively, a contradiction with the definition of I' . Q.E.D.

Proof of Proposition 5. Since total congestion can get large and there is vanishing incremental congestion, Proposition 2 and Theorem 2 apply to deduce complete learning. Given complete learning, it is straightforward that vanishing incremental congestion implies that asymptotic net congestion must equal $1/k$ (resp. $-1/k$) a.s. when the state is $\theta = 1$ (resp. $\theta = -1$).³¹ From the hypothesis that $\Delta(a^i) = g(\rho(a^i))$ for some continuous and injective function $g : [0, 1] \rightarrow \mathbb{R}$, it follows that whenever Δ converges, so must the action frequency, ρ . Plainly, if $\Delta \rightarrow 1/k$, then $\rho \rightarrow g^{-1}(1/k)$ whereas if $\Delta \rightarrow -1/k$ then $\rho \rightarrow g^{-1}(-1/k)$. Noting that the fraction of agents choosing the superior action in state -1 is $1 - \rho$ and using the neutral prior yields the expression stated in the proposition. Q.E.D.

Proof of Theorem 4. As portions of this proof proceed quite similarly to the proof of Theorem 1, we only sketch the argument.

First, note that instead of depending upon only one posterior threshold (the p^* of Lemma 1), individual decision-making now depends upon two posterior thresholds, denoted by $p^H \geq p^L$: an agent chooses $a = 1$ if $p > p^H$; $a = 0$ if $p \in (p^L, p^H)$; and $a = -1$ if $p < p^L$. This threshold structure owes to the state-independent payoff of the outside option. Consequently, in addition to the $\bar{\Delta}$ and $\underline{\Delta}$ functions, we will also have for each $a \in \{-1, 1\}$, two additional threshold functions, say $\bar{c}^a(l; k)$ and $\underline{c}^a(l; k)$, satisfying $\bar{c}^a(\cdot) \geq \underline{c}^a(\cdot)$ such that $c((a^{i-1}, a)) > \bar{c}^a$ (resp. $< \underline{c}^a$) implies $a_i \neq a$ (resp. $a_i \neq 0$) because the outside option is preferred to a (resp. is worse than a) action a no matter i 's private belief.³²

Now, towards a contradiction, fix a public likelihood ratio $x > 0$ such that for all $\varepsilon > 0$, $l_\infty^k \in B_\varepsilon(x)$ with positive probability for a sequence of $k \rightarrow 0$, just as in the proof of Theorem 1. Consider any $a_i \in \{-1, 1\}$. Just like in the proof of Theorem 1, we can show that that $l_i^k, l_{i+1}^k \in B_\varepsilon(x)$ implies that $c((a^{i-1}, a_i)) \notin [c_*(a_i, x, \varepsilon; k), c^*(a_i, x, \varepsilon; k)]$, where $\bar{c}^a(x; k) > c^*(a_i, x, \varepsilon; k) > c_*(a_i, x, \varepsilon; k) > \underline{c}^a(x; k)$. Analogous to (5) in the proof of Theorem 1, the ‘‘exclusion interval’’ $c^*(a_i, x, \varepsilon; k) - c_*(a_i, x, \varepsilon; k) \rightarrow \infty$ as $k \rightarrow 0$. Note that this is in addition to retaining the $[\Delta_*(x, \varepsilon; k), \Delta^*(x, \varepsilon; k)]$ exclusion interval as in the proof of Theorem 1; the point is that now, a_i not only reveals that the other uncertain-payoff action is undesirable to i , but moreover that the outside option is also undesirable.

³¹Note that while complete learning would be implied by large total congestion and gaps (which is weaker than vanishing incremental congestion), gaps is not sufficient to deduce that asymptotic net congestion must converge a.s.

³²For example, $\bar{c}^1(l; k) = p(\bar{b}, l)/k$ and $\underline{c}^1(l; k) = p(\underline{b}, l)/k$. Note also that the interpretation of the $\underline{\Delta}$ and $\bar{\Delta}$ functions is now different than without the outside option: for example, $\Delta \geq \bar{\Delta}$ does not imply that an agent necessarily takes action -1 , rather, it only implies that the agent will not take action 1 .

On any sample path, unbounded total congestion and full decay imply that at least two actions in $\{-1, 0, 1\}$ must be played infinitely often. Given small ε , let t be a time on some sample path such that $l_i^k \in B_\varepsilon(x)$ for all $i \geq t$. There are two exhaustive possibilities:

1. If there is an infinite number of consecutive pairs of agents such that the first agent takes action 1 and the second action -1 , then the argument in the proof of [Theorem 1](#) applies due to the existence of the set I described there.
2. If the above fails, then for some $a \neq 0$, there is an infinite number of consecutive actions $(0, a)$.³³ Now we just apply the same argument as in the proof of [Theorem 1](#) to c^a rather than Δ : i.e., $c((a^i, a)) \geq c^*(\cdot)$ and $c((a^i, 0, a)) \leq c_*(\cdot)$, but since we established above $c^*(a, \cdot) - c_*(a, \cdot) \rightarrow \infty$ as $k \rightarrow 0$, this contradicts bounded increments. *Q.E.D.*

³³To confirm this, suppose there is only a finite number of consecutive action pairs $(0, -1)$ and $(0, 1)$. Then there is only a finite number of action 0's; hence, given that two actions are played infinitely often, the first case applies.

References

- Acemoglu, Daron, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar**, “Bayesian Learning in Social Networks,” *Review of Economic Studies*, 2011, 78 (4), 1201–1236.
- Aksin, Zeynep, Baris Ata, Seyed Morteza Emadi, and Che-Lin Su**, “Structural Estimation of Callers’ Delay Sensitivity in Call Centers,” *Management Science*, December 2013, 59 (12), 2727–2746.
- Ali, S. Nageeb and Navin Kartik**, “Herding with Collective Preferences,” *Economic Theory*, November 2012, 51 (3), 601–626.
- Avery, Christopher and Peter Zemsky**, “Multidimensional Uncertainty and Herd Behavior in Financial Markets,” *American Economic Review*, 1998, 88 (4), 724–748.
- Banerjee, Abhijit**, “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 1992, 107 (3), 797–817.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch**, “A Theory of Fads, Fashion, Custom and Cultural Change as Information Cascades,” *Journal of Political Economy*, 1992, 100 (5), 992–1026.
- Callander, Steven**, “Bandwagons and Momentum in Sequential Voting,” *Review of Economic Studies*, 2007, 74, 653–684.
- Choi, Jay Pil**, “Herd Behavior, the “Penguin Effect,” and the Suppression of Informational Diffusion: an Analysis of Informational Externalities and Payoff Interdependency,” *Rand Journal of Economics*, 1997, 28 (3), 407–425.
- Cipriani, Marco and Antonio Guarino**, “Herd Behavior and Contagion in Financial Markets,” *The B.E. Journal of Theoretical Economics*, 2008, 8 (1).
- Cripps, Martin and Caroline Thomas**, “Strategic Experimentation in Queues,” February 2014. unpublished.
- Dasgupta, Amil**, “Social Learning with Payoff Complementarities,” May 2000. London School of Economics.
- **and Andrea Prat**, “Information aggregation in financial markets with career concerns,” *Journal of Economic Theory*, 2008, 143, 83–113.
- Debo, Laurens G., Christine Parlour, and Rajan Uday**, “Signaling Quality via Queues,” *Management Science*, May 2012, 58, 876–891.
- Dekel, E. and M. Piccione**, “Sequential Voting Procedures in Symmetric Binary Elections,” *Journal of Political Economy*, 2000, 108, 34–55.
- Drehmann, Matthia, Jörg Oechssler, and Andreas Roeder**, “Herding with and without payoff externalities – an internet experiment,” *International Journal of Industrial Organization*, April 2007, 25 (2), 391–415.
- Gaigl, Alexander**, “Observational Learning & Strategic Externalities,” December 2009. mimeo, University of Mannheim.

- Glosten, Lawrence R. and Paul Milgrom**, “Bid, ask and transaction prices in a market-maker market with heterogeneously informed traders,” *Journal of Financial Economics*, 1985, *14* (1), 71–100.
- Goeree, Jakob, Thomas Palfrey, and Brian Rogers**, “Social Learning with Private and Common Values,” *Economic Theory*, 2006, *28* (2), 245–264.
- Hassin, Refael and Moshe Haviv**, *To Queue or not to Queue: Equilibrium Behavior in Queueing Systems*, Boston: Kluwer Academic Publishers, 2003.
- Herrera, Helios and Johannes Hörner**, “A Necessary and Sufficient Condition for Information Cascades,” 2012. mimeo, Columbia University and Yale University.
- Koessler, Frédéric, Charles Noussair, and Anthony Ziegelmeyer**, “Parimutuel betting under asymmetric information,” *Journal of Mathematical Economics*, July 2008, *44* (7–8), 733–744.
- Lee, In-Ho**, “On the Convergence of Informational Cascades,” *Journal of Economic Theory*, 1993, *61*, 395–411.
- Lobel, Ilan, Daron Acemoglu, Munther A. Dahleh, and Asuman Ozdaglar**, “Rate of Convergence of Learning in Social Networks,” *Proceedings of the American Control Conference*, 2009.
- Lu, Yina, Andrés Musalem, Marcelo Olivares, and Ariel Schilkrut**, “Measuring the Effect of Queues on Customer Purchases,” *Management Science*, 2013, *59* (8), 1743–1763.
- Naor, P.**, “The Regulation of Queue Size by Levying Tolls,” *Econometrica*, January 1969, *37* (1), 15–24.
- Neeman, Zvika and Gerhard Orosel**, “Herding and the winner’s curse in markets with sequential bids,” *Journal of Economic Theory*, 1999, *85*, 91–121.
- Owens, David**, “An Experimental Analysis of Observational Learning with Payoff Externalities,” 2010. mimeo.
- Park, Andreas and Hamid Sabourian**, “Herding and Contrarian Behavior in Financial Markets,” *Econometrica*, 2011, *79* (4), 973–1026.
- Smith, Lones and Peter Sørensen**, “Pathological Outcomes of Observational Learning,” *Econometrica*, 2000, *68* (2), 371–398.
- and —, “Rational Social Learning by Random Sampling,” Mimeo 2008.
- Veeraraghavan, Senthil K. and Laurens G. Debo**, “Herding in Queues with Waiting costs: Rationality and Regret,” *Manufacturing & Service Operations Management*, 2011, *13* (3), 329–346.