Columbia University

Department of Economics
Discussion Paper Series

Congested Observational Learning

Eric Eyster, Andrea Galeotti, Navin Kartik, Matthew Rabin

Discussion Paper No.: 1213-07

Department of Economics
Columbia University
New York, NY 10027

November 2012

Congested Observational Learning

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

August 7, 2012

Abstract

We study observational learning in environments with congestion costs: as more of one's predecessors choose an action, the payoff from choosing that action decreases. Herds cannot occur if congestion on an action can get so large that an agent would prefer to take a different action no matter his beliefs about the state. To the extent that "switching" away from the more popular action also reveals some private information, social learning is improved. The absence of herding does not guarantee complete learning, however, as information cascades can occur through perpetual but uninformative switching between actions. Our main contribution is to provide conditions on the nature of congestion costs that guarantee complete learning and conditions that guarantee bounded learning. We also show that congestion costs have ambiguous effects on the proportion of agents who choose the superior action. 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 Dessein for his discussion at the 2012 Utah Winter Business Economics Conference, and audiences at the LSE, University of Essex, and University of Vienna for their helpful comments. Andrea Galeotti is grateful to the European Research Council for support through ERC-starting grant (award no. 283454).

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 economic contexts: as more individuals purchase a product, prices may increase (Avery and Zemsky, 1998), waiting times for made-to-order goods such as airplanes may lengthen, short-term supplies may run out, or quality of service may worsen. Alternatively, they can arise from a taste for "anti-conformity": people may have an intrinsic preference for avoiding options to which others flock.

Our model in Subsection 2.1 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 or B. One of the actions is superior to the other; while all agents share common preferences on this dimension, each 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 turns out to be the more interesting case, but otherwise make no assumptions about the distributions of private signals.

We enrich the standard model by assuming that in addition to desiring the superior action, agents may also dislike taking an action more when more of their predecessors have chosen it. We parameterize how much agents care about congestion costs relative to taking the superior action by a marginal-rate-of-substitution parameter $k \geq 0$. When $k = 0$, the model collapses to the standard one without congestion costs. Our main interest is in the long-run outcomes of games where $k > 0$: 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? We also study conditions under which a herd — when all subsequent agents take the same action — and/or an information cascade — when all subsequent agents ignore their private information — can and cannot arise.

Section 3 develops some preliminaries about individual decision-making as a function of private beliefs and the public history of actions. We then study the asymptotic properties of social 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.[1] Moreover, learning is bounded in the sense that society never learns with certainty, even asymptotically, which action is superior. We find that 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 then the other) that an agent may face is bounded above in absolute value by some threshold that is sufficiently small relative to the marginal-rate-of-substitution parameter k.

By contrast, if the net congestion cost can become sufficiently high, then it is clear that

[1] Throughout this introduction, we suppress technical details such as "almost sure" caveats.

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. When an agent's "switch" from A to B also depends on his private information, social learning is enhanced. However, the impossibility of herding does not suffice to guarantee complete asymptotic learning: after some time, agents may 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 k. Indeed, there are interesting classes of congestion costs for which, no matter the value of k, such an outcome will necessarily arise in finite time.

We show that whether such cyclical behavior and cascades can occur depends on the incremental effect that any one agent's action has on the net congestion cost faced by his successors. In particular, we identify properties of congestion costs that ensure bounded learning asymptotically even when herds cannot arise. Conversely, we provide conditions that guarantee complete asymptotic learning. It is worth noting that 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 total congestion costs are bounded there is bounded asymptotic learning once k is small enough. However, when total congestion can grow arbitrarily large as sufficiently many consecutive agents play the same action, then, under a mild additional condition, there is essentially complete asymptotic learning as $k \to 0$. Intuitively, only at near-certain public beliefs can vanishingly small incremental congestion costs produce the sort of uninformative cycling that stalls learning. In a sense, this result can be interpreted as a fragility of the conventional bounded-learning result; Section 4 clarifies to what extent such an interpretation is valid.

Section 5 turns to the asymptotic properties of the frequency of actions. We first provide suitable conditions on congestion costs under which the action frequency would converge if the superior action were known to all agents from the outset. It turns out that under these conditions the action frequency in the game of observational learning also converges, and moreover, converges to exactly the same value as if the superior action were known all along.[2]

Using these results, we show that congestion costs have ambiguous effects on the proportion of agents taking the superior action in the long-run. Although the presence of congestion costs may improve society's learning about which action is superior, agents may end up choosing the superior action less frequently in expectation than they would in the absence of congestion costs, and even under autarky. However, in other cases, eventually a larger fraction of agents take the superior action under congestion costs than without congestion costs. Indeed, the

[2]This occurs in some cases due to complete learning but in other cases despite bounded learning.

fraction of agents who choose the superior action under congestion costs may even converge to one. As a corollary, we propose an extremely simple tax scheme (whose proceeds can be redistributed, if desired) that a social planner can use in the model without congestion to achieve the first-best outcome in the long run.

In Section 6, we discuss two economic applications that fit into our general framework. In the first application, congestion costs are effectively induced by how market prices evolve over time. Our analysis 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. The second application explores a queuing model where players 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.

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 2.2), where the congestion cost of an action is a linear function of the number of predecessors who have taken that action.[3] 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 accomating 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 do not hold more generally but matter crucially for conclusions about learning, information cascades, and herds. For example, complete asymptotic 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, Parlour and Uday (2012) develop queuing models where agents observe only an aggregate statistic of predecessors' choices but not the entire history. Because Bayesian inference in this setting is extremely complex, these papers do not analyse asymptotic learning but instead characterise some properties of equilibrium play in early rounds. Drehmann et al. (2007) conduct an 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.[4]

Our work also relates to Avery and Zemsky (1998), who build on Glosten and Milgrom (1985)'s model of sequential trade for an asset of common but unknown value. In Avery and Zemsky's simplest variant, a market-maker sets the price of the risky asset at the start of every period t to equal the public belief based on all information revealed through period

[3]We learned of Gaigl's work only after circulating a prior draft of this paper.

[4]Owens (2010) also presents an experiment on observational learning with payoff externalities, finding that decisions are highly responsive to both positive and negative payoff externalities.

$t-1$.[5] Because the price fully incorporates all earlier traders' private information, each trader buys when her private information about the value of the asset is positive and sells when negative. Because each trader acts informatively, the market price converges to the asset's true value. Market prices in Avery and Zemsky (1998) play a similar role to congestion costs in our model: holding 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 generalising their model 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 players act uninformatively, and on the other hand that different mechanisms of price formation can substantially alter the conclusion of complete learning.

We note that are models of observational learning without congestion costs in which complete learning obtains. Lee (1993) derives such a result when players' action spaces are a continuum—rich enough to reveal their posteriors—and preferences satisfy some reasonable properties. Even with only a finite number actions, Smith and Sørensen (2000) show that when players' private beliefs are unbounded, complete learning obtains when players have the same preferences.[6] By contrast, complete learning can obtain in our model even under a binary action space and bounded private beliefs. Furthermore, in the settings explored by Lee (1993) and Smith and Sørensen (2000), *every* player's action is informative, i.e. depends upon his private signal; as already mentioned, this is typically not the case here.

Finally, this paper contributes to an expanding 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, 2011, 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.[7]

[5] Thus, in the simplest variant of Avery and Zemsky (1998), the market-maker loses money on average. Their richer model with noise traders does not share this feature.

[6] Intuitively, information cascades cannot occur under unbounded private beliefs: whatever the current public belief, a player can receive an opposing signal strong enough to overturn it. This is an incomplete intuition because the absence of information cascades is compatible with a failure of complete learning, but this requires bounded private beliefs (Herrera and Hörner, 2011). Under bounded private beliefs, Goeree et al. (2006) prove a learning result when players' preferences include a full-support private-values component in addition to the common objective of matching action to state; see also Goeree et al. (2007). Acemoglu, Dahleh, Lobel and Ozdaglar (2011) show that learning can occur under bounded private beliefs if not all players necessarily observe all predecessors' choices.

[7] Drehmann et al. (2007) include treatments with forward-looking payoff externalities. They find that subjects' behavior in these treatments do not differ significantly from myopic behavior and suggest that a purely backward-looking analysis might make reasonably good predictions even in settings with forward-looking considerations.

2 Congestion Costs

2.1 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 players take actions in sequence, each observing the entire history of actions. Before choosing an action, player 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 a player's belief that $\theta = 1$ after observing her signal but ignoring the history of play, computed by Bayes rule using the uniform prior. Denote the private belief of player i as $p_i \in (0, 1)$. Given the state $\theta \in \{-1, 1\}$, the private-belief stochastic process $\langle p_i \rangle$ is conditionally i.i.d. with conditional c.d.f. $F^{(\theta)}$. We assume that no private signal perfectly reveals the state of the world, which implies that $F^{(1)}$ and $F^{(-1)}$ are mutually absolutely continuous and have common support. Denote the convex hull of that support by $[\underline{b}, \overline{b}] \subseteq [0, 1]$. To avoid trivialities, signals must be informative, which implies that $\underline{b} < 1/2 < \overline{b}$. We focus in this paper on *bounded private beliefs*: $\underline{b} > 0$ and $\overline{b} < 1$.[8] Notice that this setting allows for continuous or discrete signals.

Denote each player's action by $a_i \in \{-1, 1\}$ and let $a^i := (a_1, \ldots, a_i)$ denote a history. Player i's preferences are given by a von-Neumann-Morgenstern utility function

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

where $\mathbb{1}_{\{\cdot\}}$ denotes an 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 a player's index, different players may be affected differently in terms of congestion by common predecessors, and furthermore, different players 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 $k \to 0$.

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

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

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

[8] The case of unbounded private beliefs ($\underline{b} = 0$ and $\overline{b} = 1$) is briefly discussed in the conclusion.

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

Our other maintained assumption places a mild bound on the extent to which players can affect their immediate successors' net congestion cost:

Assumption 2 (Quasi-bounded increments)**.** *For all $m \in \mathbb{R}$, there exists $m' \in [m,\infty)$ and $m'' \in (-\infty, m]$ such that $\Delta(a^i) \leq m \implies \Delta((a^i,1)) \leq m'$ and $\Delta(a^i) \geq m \implies \Delta((a^i,-1)) \geq m''$.*

Assumption 2 is related to, but weaker than, requiring the incremental effect of an action on net congestion cost to be bounded. It is weaker because it doesn't impose a bound on the incremental effects of player i's action when the net congestion cost faced by i gets arbitrarily large or small. Accordingly, we refer to Assumption 2 as a "quasi-bounded increments" assumption. See fn. 10 below for an example of how this assumption can be violated.

2.2 Leading examples

It is useful to introduce two leading examples that satisfy all the maintained assumptions. In both, people 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 $i \geq 2$,

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

If we denote $y(a^i)$ as the fraction of agents $1, \ldots, i$ who have chosen $a = 1$ under a^i, then for each $i \geq 2$, $\Delta(a^i) = 2y(a^i) - 1$. More generally, instead of $\Delta(a^i)$ being linear in $y(a^i)$, we could have $\Delta(a^i) = f(y(a^i))$ for some function $f : (0,1) \to \mathbb{R}$ that is strictly increasing but otherwise arbitrary; this defines the general *proportional-costs model*.[9] ■

Example 2. In the *linear-absolute-cost model*, let

$$\eta(a^i) = \sum_{j=1}^{i} \mathbb{1}_{\{a_j = 1\}},$$

and define $c(a^1) = 0$ and for each $i \geq 2$,

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

[9] Note 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 players' choices to be $a_1 = 1$ and $a_2 = -1$.

Here, congestion depends upon the number of agents who have chosen action 1 rather than the fraction, so that $\Delta(a^i) = 2\eta(a^i) - i$.[10] A general *absolute-costs model* has $\Delta(a^i) = f(\eta(a^i)) - f(i - \eta(a^i))$ for some strictly increasing $f : \mathbb{R}_+ \to \mathbb{R}$. ■

3 Individual Decision-Making

Player i's decision depends upon her beliefs about which action is superior as well as upon net congestion costs. Let p be player 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 player i makes from the history a^{i-1}. Given q and p, i's posterior belief that the state is $\theta = 1$ is given by

$$r(p, q) = \frac{pq}{pq + (1-p)(1-q)}.$$

Let $l = (1-q)/q$ be the *public likelihood ratio* (LR), which is the inverse of the relative likelihood of state 1; low l means that is more likely that $\theta = 1$. We can rewrite $r(p, q)$ as a function of l as follows:

$$r(p, l) = \frac{p}{p + (1-p)l}. \tag{1}$$

Clearly, $r(p, l)$ is strictly increasing in p and strictly decreasing in l. The following lemma describes how posterior beliefs determine action choice.[11]

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

Proof. It suffices to compute

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

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

Using Lemma 1, we can derive the net congestion cost that renders a player indifferent between action 1 and action -1 given likelihood ratio of l and the private belief most favorable

[10] The following variant of the linear-absolute-cost example violates Assumption 2 yet satisfies Assumption 1:

$$\Delta(a^i) = \begin{cases} \eta(a^i) & \text{if } \eta(a^i) \geq \frac{i}{2} \\ \eta(a^i) - i & \text{if } \eta(a^i) < \frac{i}{2}. \end{cases}$$

To see why, consider an odd i and let a^i be such that $\eta(a^i) = i - \eta(a^i) - 1$. Then $\Delta(a^i) = \eta(a^i) - i < 0$ while $\Delta((a^i, 1)) = \eta(a^i) + 1$, which gets arbitrarily large as $i \to \infty$.

[11] Throughout the paper, we do not carefully specify how to break indifference: with continuous signals, the choice doesn't matter, and with discrete signals, it is generically irrelevant.

to action 1. Formally, for any $l \in \mathbb{R}_+$, define $\overline{\Delta}(l;k)$ to be the unique solution to $r(\overline{b}, l) = r^*(\overline{\Delta}(l;k);k)$.[12] Given a net congestion cost of $\overline{\Delta}(l;k)$ and public likelihood of l, any player—no matter what her private information—(weakly) prefers choosing $a = -1$ to $a = 1$. Note that for any p, l, and $\Delta \geq \overline{\Delta}(l;k)$, we have that $r(p,l) \leq r^*(\Delta;k)$, which, by Lemma 1, implies that the agent will uninformatively choose $a = -1$.

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

We call a player's action *informative* if it depends non-trivially upon her private beliefs, or equivalently, if $\Delta \in (\underline{\Delta}(l;k), \overline{\Delta}(l;k))$. The next lemma provides a number of useful 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 $l \to \infty$; and hence, (iii) $\overline{\Delta}(l;k) - \underline{\Delta}(l;k) \to 0$ as $l \to \infty$ or as $l \to 0$.*
3. *For any $l > 0$, $\overline{\Delta}(l;k) - \underline{\Delta}(l;k) \to \infty$ as $k \to 0$.*
4. *$\overline{\Delta}(l;\cdot)$ is strictly decreasing (resp. increasing) for any $l > 0$ that is strictly smaller (resp. larger) than $\frac{\overline{b}}{1-\overline{b}}$; 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. Equation (1) and the definitions of $\overline{\Delta}(l;k)$ and $\underline{\Delta}(l;k)$ yield

$$\begin{aligned} \overline{\Delta}(l;k) &= \frac{\overline{b} - (1-\overline{b})l}{k(\overline{b} + (1-\overline{b})l)}, \\ \underline{\Delta}(l;k) &= \frac{\underline{b} - (1-\underline{b})l}{k(\underline{b} + (1-\underline{b})l)}. \end{aligned}$$

The lemma's first two parts and the fourth part are straightforward to verify from the above formulae. For the third part, observe that by the definitions, for any $l > 0$ and $k > 0$,

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

Since the right hand side above is strictly positive, the result follows. *Q.E.D.*

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

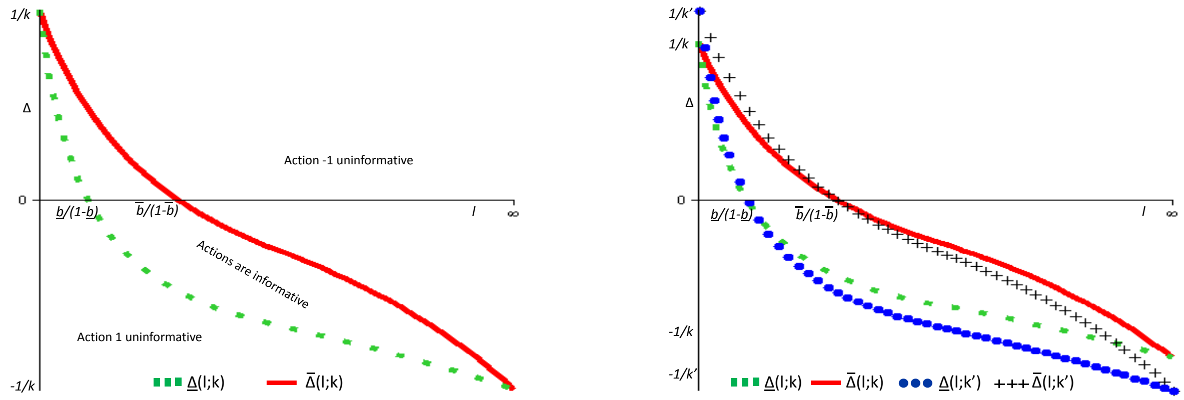

(a) Parts 1 and 2 of Lemma 2

(b) Parts 3 and 4 of Lemma 2, with $k > k'$.

Figure 1 – Net congestion threshold functions.

Figure 1(a) illustrates Lemma 2's first two observations. For net congestion costs above $\overline{\Delta}(l;k)$ (the bold line), players choose action -1 regardless of their private signal; for net congestion costs below $\underline{\Delta}(l;k)$ (the dotted line), players choose action 1 regardless of their private signal; for net congestion costs that lie between the two lines, players choose actions that depend upon their private signals. Figure 1(b) illustrates how a change in k affects both $\underline{\Delta}(l;k)$ and $\overline{\Delta}(l;k)$ (parts 3 and 4 of Lemma 2). For a given l, when k decreases, $\overline{\Delta}(l;k)$ rotates around $\overline{b}/(1-\overline{b})$: it increases when the likelihood ratio is below $\overline{b}/(1-\overline{b})$ 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, and moreover, this difference gets arbitrarily large as k vanishes.

4 Learning

4.1 Concepts

Without loss of generality, assume that 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 $\mathrm{Supp}[l_\infty^k] \subseteq [0, \infty)$.

We say that there is *complete learning* if beliefs a.s. converge to the truth, i.e. $\Pr(l_\infty^k = 0) = 1$. We say there is *incomplete learning* when $\Pr(l_\infty^k = 0) = 0$, i.e. beliefs a.s. do not converge to the truth. If there exists $\varepsilon > 0$ such that $\Pr(l_\infty^k > \varepsilon) = 1$, we say there is *bounded learning* because beliefs a.s. are bounded away from the truth. Clearly, bounded learning implies incomplete learning but not vice-versa. There is a *herd* on a sample path if after some (finite) time, all subsequent players choose the same action. There is an *information cascade* on a sample path if after some (finite) time, no player'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 (and hence incomplete) learning. There is almost surely a herd; 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]

Note that in the standard model without congestion, an information cascade implies a herd. Much of the analysis builds upon the fact that in the current model this need not be the case.

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

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 $k \to 0$. 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, we provide 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.

For a given $k > 0$, there are two potential reasons why learning may not occur. The first is standard: it is possible that players will eventually herd (and do so uninformatively). Whether a herd can occur in the present context depends on the cumulative effect of players' actions on net congestion costs. During a putative herd on action 1, say, players' 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 during an uninformative herd), the herd persists. Figure 2(a)

[13]Regarding 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 $r(\underline{b}, \underline{l}) = 1/2$. 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 $\hat{l} := \frac{1-r(\bar{b},\underline{l})}{r(\bar{b},\underline{l})}$, i.e. $\hat{l}$ is the posterior LR obtained from a public LR 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 (2011).

illustrates this possibility of uninformative herding. In contrast, if the net congestion cost eventually becomes sufficiently high, then at some point a player will necessarily choose action -1, thereby breaking the putative herd.[14]

The impossibility of uninformative herding is not *per se* sufficient for learning. The second reason why learning can fail 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 his predecessor. Naturally, such an outcome is only possible under suitable net congestion cost functions; Theorem 3 below identifies a sufficient condition. In particular, if the incremental effect that any player's action eventually has on net congestion costs becomes negligible, perpetual uninformative oscillation can never occur, no matter the value of k. This is illustrated in Figure 2(c); Theorem 2 below develops the point into a complete-learning result. Even when uninformative oscillations between actions can occur for a given value of k, they can only do so when a player's incremental effect on net congestion cost is sufficiently large relative to the public likelihood ratio. For example, in Figure 2(b), oscillation between the two "white" dots is uninformative under k, yet for $k' < k$, behavior at the lower of those points is informative because that point now lies in between the threshold curves corresponding to k'. As long as incremental congestion effects are bounded, when k becomes small uninformative oscillation cannot occur forever except at very extreme public likelihood ratios; Theorem 1 below elaborates this idea.

To further substantiate these notions, consider the linear versions of Example 1 and Example 2. In the linear-proportional-cost model, the incremental effect on net congestion cost of a player's action eventually vanishes. Hence, for any $k > 0$, perpetual uninformative oscillation cannot occur; the only force that can prevent learning is herding. When $k < 1$, congestion costs never rise to a level sufficient to break a herd. When $k \geq 1$, congestion costs eventually do become high enough to prevent herds. In the linear-absolute-cost model, congestion costs can always get high enough to prevent a herd. However, since each player's action affects net congestion cost identically, players eventually oscillate uninformatively between the two actions indefinitely. Yet as k becomes arbitrarily small, these uninformative oscillations can only occur at extreme public likelihood ratios.

The following proposition, whose proof is omitted because it follows from the general theorems in Subsections 4.3–4.5, develops these results:

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) incomplete learning for any $k > 0$; but (ii) learning with high probability in the no-congestion limit.*

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

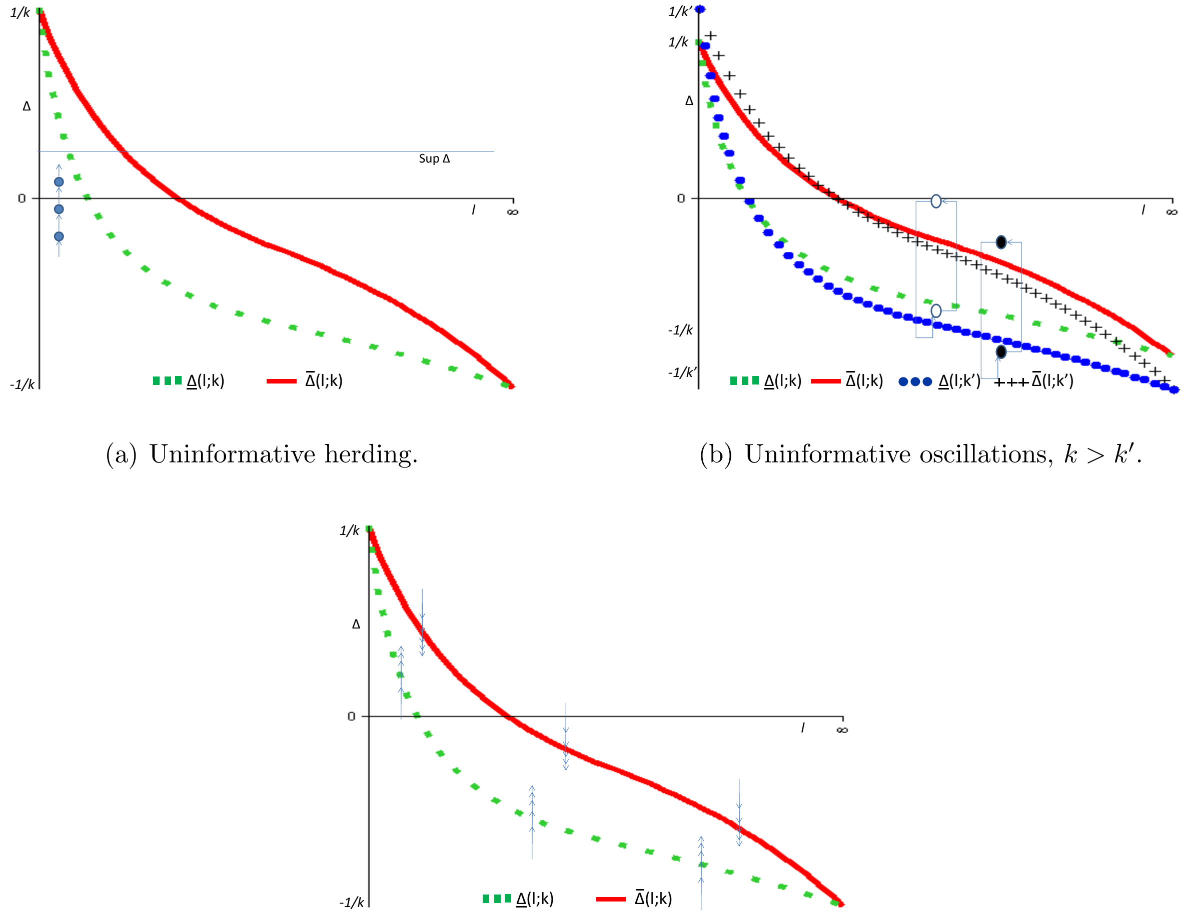

(a) Uninformative herding.

(b) Uninformative oscillations, $k > k'$.

(c) Absence of perpetual uninformative oscillations.

Figure 2 – Net-congestion costs and learning.

3. *In a non-linear proportional-costs model with an $f(y(a^i))$ function whose range is $(-\infty, \infty)$, there is complete learning for any $k > 0$ and, hence, complete learning in the no-congestion limit.*

4. *In a non-linear absolute-cost 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.*

Furthermore, one can also deduce properties about the asymptotic fraction of agents taking a particular action in the leading examples; this is studied in general in Section 5.

4.3 The no-congestion limit

We begin the formal analysis by studying learning in the no-congestion limit.

Definition 1. Total congestion is *bounded* if $-\infty < \inf \Delta(a^i)$ and $\sup \Delta(a^i) < +\infty$. Total congestion is *unbounded* if

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

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. 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.

Definition 2. There is *bounded incremental congestion* if

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

This condition simply requires that the marginal impact of any individual's choice on net congestion cost be bounded above; although it is strictly stronger than Assumption 2, it is generally a mild requirement.

Theorem 1. *If total congestion is unbounded and there is bounded incremental congestion, 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 2 (see Corollary 1); we prove here the first statement of the theorem.

Step 1: We first claim that for any $x > 0$, $x \notin \mathrm{Supp}[l_\infty^k]$ once k is small enough. Suppose not, per contra, for some $x > 0$. Then, for any $\varepsilon > 0$, there is a sequence of $k \to 0$ such that for each k, there is positive probability that $l_\infty^k \in B_\varepsilon(x)$. For $\varepsilon > 0$ small enough, if l_i^k and l_{i+1}^k are both in $B_\varepsilon(x)$, then i's private belief threshold (recall Lemma 1), call it $\hat{b}_i$, cannot lie within some interval $[b^*(x,\varepsilon), b_*(x,\varepsilon)] \subsetneq [\underline{b}, \overline{b}]$; this follows from the fact that $x > 0$ and l_{i+1}^k is derived through Bayesian-updating of l_i^k using either $b_i > \hat{b}_i$ (if $a_i = 1$) or $b_i < \hat{b}_i$ (if $a_i = -1$). Moreover,

$$\text{as } \varepsilon \to 0,\ b^*(x,\varepsilon) \to \overline{b} > 1/2 \text{ and } b_*(x,\varepsilon) \to \underline{b} < 1/2. \tag{2}$$

Just as in the proof of Lemma 2, the belief threshold $b_*(x,\varepsilon)$ can be mapped into a net-congestion threshold, $\Delta_*(x,\varepsilon;k) > \underline{\Delta}(x;k)$ through the equation

$$1 + k\Delta_*(x,\varepsilon;k) = 2b^*(x,\varepsilon), \tag{3}$$

and analogously $b^*(x,\varepsilon)$ maps into a threshold $\Delta^*(x,\varepsilon;k) < \overline{\Delta}(x;k)$. For l_i^k and l_{i+1}^k to both lie in $B_\varepsilon(x)$, it must be that $\Delta(a^{i-1}) \notin [\Delta_*(x,\varepsilon;k), \Delta^*(x,\varepsilon;k)]$. From (2) and (3), it follows that for any small enough $\varepsilon > 0$,

$$\text{as } k \to 0,\ \Delta_*(x,\varepsilon;k) \to -\infty \text{ and } \Delta^*(x,\varepsilon;k) \to +\infty. \tag{4}$$

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 $i \geq t$. 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, $I \subseteq \{t, t+1 \ldots, \}$ such that for any $i \in I$, $\Delta\left(a^i\right) \leq \Delta_*(x,\varepsilon;k)$ whereas $\Delta((a^i,1)) \geq \Delta^*(x,\varepsilon;k)$. However, (4) implies that for any $\varepsilon > 0$ small enough, $\Delta^*(x,\varepsilon;k) - \Delta_*(x,\varepsilon;k) \to \infty$ as $k \to 0$. The hypothesis of bounded incremental congestion now implies that once k is small enough, if $\Delta\left(a^i\right) \leq \Delta_*(x,\varepsilon;k)$ then $\Delta((a^i,1)) < \Delta^*(x,\varepsilon;k)$, a contradiction.

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 $\operatorname{Supp}[l_\infty^k] \subseteq [0,\varepsilon] \cup [L,\infty)$ once k is small enough.

For any $\varepsilon > 0$, by applying the above claim to a sequence of $\delta \to 0$, it holds that $\lim_{k\to 0} \Pr[l_\infty^k \leq \varepsilon] = 1$; 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. One can interpret the theorem as identifying a sense in which the bounded learning conclusion of the standard model without congestion (recall Remark 1) is fragile, but this requires recognizing in what sense our model does and does not converge to the benchmark model as $k \to 0$. 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 $k \to 0$, our model converges pointwise across players to the model without congestion. However, when total congestion is unbounded, the convergence is not uniform: the values of $\delta(i,\varepsilon)$ cannot be chosen independently of i. By contrast, there is uniform convergence when total congestion is bounded.

4.4 Complete learning for any $k > 0$

We turn to the question of what properties of congestion costs assure complete learning for an arbitrary $k > 0$. As suggested by the discussion in Subsection 4.2, one key condition is that total congestion cost can become large enough in magnitude to prevent herds. In addition, one must also ensure that players cannot perpetually oscillate between actions without conveying any information about their private beliefs.

Definition 3. For any $k > 0$, total congestion *can get large* or *is large* if

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

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 $k > 0$.

Definition 4. Congestion has *no gaps* provided that for any $\varepsilon > 0$ and any non-convergent infinite action sequence $(a_1, \ldots)$: 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 4 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.[15]

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 $\mathrm{Supp}[l_\infty^k] = \{0\}$. Suppose, per contra, that $x > 0$ and $x \in \mathrm{Supp}[l_\infty^k]$. Following the logic developed in the proof of Theorem 1 and using the notation introduced there, we conclude that for any $\varepsilon > 0$ small enough, there must be a sample path of actions $(a_1, \ldots)$ and some time t such that

$$\text{for all } i \geq t\text{: either } \Delta(a^i) \leq \Delta_*(x, \varepsilon; k) \text{ or } \Delta(a^i) \geq \Delta^*(x, \varepsilon; k). \tag{5}$$

Since large total congestion implies

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

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 quasi-bounded increments (Assumption 2), (5) 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

[15] Since 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.

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 (5). *Q.E.D.*

In terms of the no-congestion limit, Theorem 2 strengthens the positive conclusion of Theorem 1 but requires congestion to have no-gaps rather than bounded incremental congestion. More importantly, Theorem 2 can be applied to arbitrary $k > 0$.[16]

In Subsection 5.1, we introduce a condition called *vanishing incremental congestion* (see Definition 7) that often may be easily verified in applications and implies that congestion has no gaps. For example, the proportional-cost model with a continuous $f(\cdot)$ satisfies this stronger condition (see fn. 20); moreover, total congestion is large in this model if and only if, $(-1/k, 1/k) \subseteq range[f(\cdot)]$. In particular, total congestion gets large in the linear-proportional-cost model if and only if $k \geq 1$. Similarly, in the absolute-cost model total congestion is large 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 $range[f(\cdot)]$ is bounded, the absolute-cost model satisfies the no gaps condition because it satisfies the stronger property of vanishing incremental congestion.

4.5 Non-learning for any $k > 0$

Our final set of learning results derive sufficient conditions for bounded (and hence incomplete) 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 $(a_1, \ldots)$, there exists i^* such that for all $i > i^*$, $\Delta(a^i) \notin (1/k - C(k), 1/k) \cup (-1/k, -1/k + C(k))$.

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. We will show subsequently how the gaps condition subsumes the intuitions provided in Subsection 4.2 about incomplete learning when congestion costs on one action never get large enough (see Proposition 2 below) and when the incremental effect of a player's action on successors' net congestion costs is never negligible.

[16] The proof of Theorem 2 explicitly uses quasi-bounded increments (Assumption 2). To see the relevance of this assumption, consider the modified version of the linear absolute cost model introduced in fn. 10 and recall that this example violates Assumption 2. Observe that for any i and a^i, either $\Delta(a^i) \geq i/2$ or $\Delta(a^i) < -i/2$; hence given any bounded interval, S, if i is large enough then for any a^i, $\Delta(a^i) \notin S$. Consequently, congestion has no gaps and total congestion is unbounded. However, for any $k > 0$, learning is necessarily incomplete because for any $i > 2k$ and any a^i, $\Delta(a^i) > 1/k$ or $\Delta(a^i) < -1/k$ and hence i's action is uninformative. In other words, in any sample path of actions, an information cascade arises no later than time $2k + 1$.

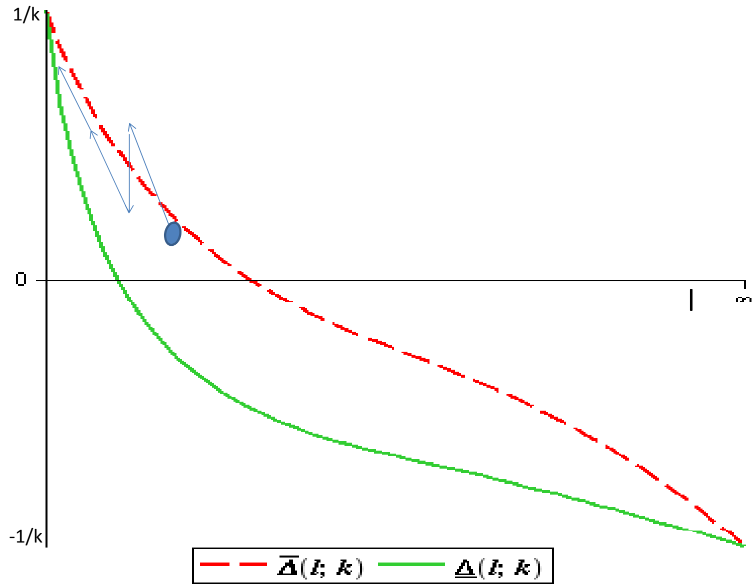

Figure 3 – A sample path along which there is learning.

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

$$\text{the range of } \Delta(\cdot) \text{ has no finite limit point.}^{17} \tag{6}$$

Clearly, the linear absolute cost model satisfies (6) because the range of $\Delta(\cdot)$ in that case is the integers.

Theorem 3. *For any $k > 0$, if congestion has gaps then there is bounded learning. Furthermore, if, for all $k > 0$ small enough, 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 $(a_1, \ldots)$. Since $1/k > \overline{\Delta}(l; k) > \underline{\Delta}(l; k)$ 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 $l_i^k < \frac{1-r(\bar{b},\hat{l})}{r(\bar{b},\hat{l})}$, where the latter term is strictly larger than 0. Since the sample path was arbitrary, it follows that there is bounded learning.

[17] For a given $k > 0$, an even weaker sufficient condition is that neither $1/k$ nor $-1/k$ be a limit point of the range of $\Delta(\cdot)$.

For the second statement, assume there is some $z > 0$ 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, $\min \operatorname{Supp}[l_\infty^k] \geq \frac{1-r(\bar{b},\hat{l})}{r(\bar{b},\hat{l})}$, 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 the same constant $C(k)$ in Definition 5 is required to apply for all infinite action sequences. If, instead, the constant could depend on the action sequence, then we would only be able to conclude incomplete learning rather than bounded learning.

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 $-1/k < \inf \Delta(a^i)$ and $\sup \Delta(a^i) < 1/k$.

Proposition 2. *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.*

Proof. 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.*

Combining Theorem 3 and Proposition 2 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 a.s., information cascades need not. Moreover, at any point where a cascade arises, so too does a herd. Hence, with congestion costs, information cascades need not usher in herds. The next result describes properties of the net congestion cost that guarantee the presence of an information cascade and simultaneously rule out herding. These conditions are satisfied, for example, by the linear-absolute-cost model.

Proposition 3. *If total congestion is unbounded and (6) holds, then for any $k > 0$ there is almost surely an information cascade.*

Proof. Fix $k > 0$ and pick any sample path $(a_1, \ldots)$; let z be limit public likelihood ratio 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 (6), 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.*

5 Action Frequency

In this section we study the asymptotic frequency of actions under congestion costs. Under suitable conditions, we compare this frequency under incomplete information about the state with the corresponding complete-information benchmark, and also deduce the fraction of agents who asymptotically choose the superior action. We then provide an interpretation of congestion costs in term of transfers and apply our results to derive a simple transfer scheme that a social planner can use to approximate first-best welfare in the standard model without congestion costs.

5.1 Action convergence

We first focus on convergence of the action frequency, i.e. $\lim_{i\to\infty} y\left(a^i\right) = \lim_{i\to\infty} \frac{|\{a_j=1:j\le i\}|}{i}$.[18] Contrary to the standard model without congestion, this limit may not exist even if the state were known, as illustrated below. We will say that the *correct action frequency* is the limit action frequency that would obtain if the true state were known, assuming this limit exists.

Fairly strong conditions are needed to ensure the existence of a correct action frequency; neither the no-gaps condition nor the gaps condition suffice. Accordingly, we will introduce two new conditions, one stronger than no gaps and the other stronger than gaps.

Definition 7. There is *vanishing incremental congestion* if for any $\varepsilon > 0$, any infinite sequence of actions $(a_1, \ldots)$, and any interval $(\underline{s}, \overline{s})$ with $\infty > \overline{s} > \underline{s} > -\infty$, there exists an index i' such that if $i > i'$ and $\Delta(a^i) \in (\underline{s}, \overline{s})$, then $\left|\Delta(a^{i+1}) - \Delta(a^i)\right| < \varepsilon$.

Vanishing incremental congestion does not compel incremental congestion to die off if total congestion is growing unboundedly, a generality which will be useful to subsume some

[18]We have defined action frequency convergence in terms of $a = 1$; this is without loss, since this converges if and only if the frequency of $a = -1$ converges.

examples (cf. fn. 20). It requires that given any infinite action sequence, eventually any two successive players i and $i+1$ face net congestions that are arbitrarily similar.[19] For example, the proportional cost model with any continuous $f(\cdot)$ satisfies vanishing incremental congestion.[20]

Proposition 4. *Vanishing incremental congestion implies that congestion has no gaps.*

Proof. Fix any non-convergent infinite sequence of actions $(a_1, \ldots)$ and any $\varepsilon > 0$. Pick any bounded interval $S \subseteq \mathbb{R}$ and any infinite set of agents $I_S = \{n_1, n_2, \ldots\}$ 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) $x^* = \lim_{k\to\infty} \Delta(a^{n_k})$ 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 $\Delta\left(a^i\right) \in (x^* - \varepsilon/2, x^* + \varepsilon/2)$ 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 $n \in I_S$, $\Delta(x^n) \in (x - \varepsilon, x + \varepsilon)$. 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.*

Definition 8. There is *constant incremental congestion* if there exists $N \in \left\{..., \frac{1}{3}, \frac{1}{2}, 1, 2, 3, ...\right\}$ and $x \geq 0$ such that in any infinite sequence $(a_1, \ldots)$, there is some i' such that for any $i > i'$, (i) $a_i = 1 \implies \Delta(a^i) - \Delta(a^{i-1}) = x$; and (ii) $a_i = -1 \implies \Delta(a^{j-1}) - \Delta(a^j) = Nx$.

Constant incremental congestion requires that in any infinite sequence of actions eventually (after some time that may depend on the sequence), two properties must hold: (i) the incremental net congestion effect of any action 1 is constant and analogously for action -1; (ii) while possibly different from each other, these two constants must be integer multiples.[21] Note that the multiple must be independent of the particular action sequence. The linear-

[19] Vanishing incremental congestion essentially combines two ideas: first, player i's own action should not affect net congestion by much; second, and more subtle, the manner by which the actions of players $1, \ldots, i-1$ enter into player i's preferences as congestion must be similar to how they enter player $i+1$'s preferences. Although reasonable in many contexts, vanishing incremental congestion rules out various kinds of time-varying congestion costs.

[20] This is because given any interval $[y_1, y_2]$ (with $y_1 > 0$ if $\lim_{y\to 0} f(y) = -\infty$ and $y_2 < 1$ if $\lim_{y\to 1} f(y) = \infty$), any infinite sequence of actions in which the net congestion always lies within this interval has the following property: given any $\varepsilon > 0$, there is i' such that for any a^i with $i > i'$,

$$\max\{f(y((a^i, 1))) - f(y(a^i)), f(y(a^i)) - f(y((a^1, -1)))\} < \varepsilon.$$

This follows from the continuity of $f(\cdot)$ and that eventually any one player's choice has negligible effect on $y(\cdot)$.

[21] Since these properties are only required after some arbitrarily large finite time, it would be more accurate to call the condition "*eventually*-constant incremental congestion"; we drop the "eventually" qualifier to ease terminology.

absolute-cost model satisfies constant incremental congestion with $N = 1$. It is transparent that constant incremental congestion implies that congestion has gaps.

Lemma 3. *If all players know the state (i.e. in the complete-information benchmark):*

1. *if total congestion is small, the action frequency converges to* 1 *(resp.* 0*) if* $\theta = 1$ *(resp.* $\theta = -1$*);*
2. *if total congestion is large and there is constant incremental congestion, the action frequency converges to* $N/(N+1)$ *for all* $\theta \in \{-1, 1\}$*, where* N *is from the definition of constant incremental congestion;*
3. *if total congestion is large and there is vanishing incremental congestion, and in addition if* $\Delta(a^i)$ *can be written as one-to-one function of* $y(a^i)$*, then the action frequency converges to* y *such that* $\Delta(y) = 1/k$ *(resp.* $\Delta(y) = -1/k$*) if* $\theta = 1$ *(resp.* $\theta = 0$*).*

Proof. We prove the result assuming $\theta = 1$; analogous arguments apply when $\theta = 0$. The first part is trivial: if total congestion is small and $\theta = 1$ is known, then all agents choose $a = 1$. For the second part, assume constant incremental congestion and that total congestion is large. Assume $N \geq 1$; the logic is analogous if $N < 1$. On any sample path, eventually the sequence of actions looks like

$$\left(\ldots, -1, \underbrace{1, 1, \ldots, 1}_{N \text{ times}}, -1, \underbrace{1, 1, \ldots, 1}_{N \text{ times}}, -1, \ldots \right),$$

with $\Delta(a^i)$ jumping from above $1/k$ to below at each action of -1, and then staying below $1/k$ until the N-th consecutive action of 1 switches it to above $1/k$. The conclusion now follows from the fact that this eventual pattern determines the asymptotic action frequency.[22]

For the third part, assume that total congestion is large and that there is vanishing incremental congestion. Then, asymptotically $\Delta = 1/k$, and it follows that if Δ is a 1-1 function of the action frequency, then the action frequency must converge.[23] *Q.E.D.*

Example 1, continued. In the linear-proportional-cost model $\Delta(a^i) = 2y(a^i) - 1$ and hence, when $k \geq 1$, Part 3 of Lemma 3 implies that the correct action frequency is $y^* = \frac{1+k}{2k}$ if $\theta = 1$ while $y^* = \frac{k-1}{2k}$ if $\theta = -1$. On the other hand, for $k < 1$, Part 1 of the Lemma implies that $y^* = 1$ when $\theta = 1$ and $y^* = 0$ when $\theta = -1$. ■

Example 2, continued. Recall that the linear-absolute-cost model has a constant incremental congestion, with $N = 1$. It follows from Part 2 of Lemma 3 that in this case the correct action frequency is $1/2$. Note that, perhaps surprisingly, this is independent of both k and

[22] Note that we are using here the fact that the value of N is independent of the sample path of actions.

[23] It is important that Δ be a 1-1 function of the action frequency; perhaps surprisingly, the action frequency is not guaranteed to converge when total congestion can get large and there is vanishing incremental congestion. A counter-example can be constructed where the sequence of actions consists of a block of 1's followed by a block of -1's and so on, with the length of each block being such that sufficiently far in the sequence, whenever a block of 1's is over, $y(\cdot) \approx 1$, whereas whenever a block of -1's is over, $y(\cdot) \approx 0$.

the state θ. The reason is that even though k and θ have a substantial effect on the number of initial agents who choose the superior action, eventually agents will just alternate between action 1 and -1. Next, consider a generalization of the linear-absolute-cost model to one where $c\left(a^{i-1},1\right)=\delta\sum_{j=1}^{i-1}\mathbb{1}_{\{a_j=1\}}$ for some $\delta>0$ whereas $c\left(a^{i-1},-1\right)=\sum_{j=1}^{i-1}\mathbb{1}_{\{a_j=1\}}$. Then $\Delta\left(a^i,1\right)-\Delta\left(a^i\right)=\delta$ while $\Delta\left(a^i\right)-\Delta\left(a^i,-1\right)=1$. So there is still a constant incremental congestion, but now $N=1/\delta$. Therefore, the correct action frequency is $1/(1+\delta)$. ■

The next proposition provides sufficient conditions for action-frequency convergence and relates them to the correct action frequency.

Proposition 5.

1. *If total congestion is small then a herd occurs almost surely (hence, almost surely there is action convergence), and with positive probability is on the inferior action.*
2. *If total congestion can get large and there is constant incremental congestion, then the action frequency converges almost surely, and when it converges, the limit is the correct action frequency.*
3. *If total congestion can get large, there is vanishing incremental congestion, and $\Delta(a^i)$ can be written as one-to-one function of $y(a^i)$, then the action frequency does converge almost surely and the limit is the correct action frequency.*

Proof. Part 1: Since learning is bounded (by Corollary 1), the public LR either converges to some $l_\infty^k\in(0,l')$, where l' is such that $\underline{\Delta}(l',k)=\inf\Delta(\cdot)$, or to some $l_\infty^k>l''$ where l'' is such that $\underline{\Delta}(l'',k)=\sup\Delta(\cdot)$. In the former case, at some point players will start choosing action 1 forever, i.e. there is a herd on action 1. In the latter case, at some point players will start choosing action -1 forever, i.e. there is a herd on action 0. So a.s. there is a herd. Since private beliefs are bounded, both herds can occur with positive probability.

Part 2: The public LR a.s. converges to some l_∞^k, and, because of bounded learning (by Theorem 3 and that constant incremental congestion implies gaps), $l_\infty^k>0$. Since total congestion is large, it must be that asymptotically Δ is oscillating from above $\overline{\Delta}(l_\infty^k,k)$ to below $\underline{\Delta}(l_\infty^k,k)$ (not necessarily switching based on just one action in both directions, of course). Since there is constant incremental congestion with factor N, the same argument as in the proof of Lemma 3 applies.

Part 3: If total congestion can get large and there is vanishing incremental congestion, then Proposition 4 and Theorem 2 implies that there is complete learning. Hence, $l_i^k\to 0$ a.s. Moreover, by vanishing incremental congestion, we must asymptotically a.s. have $\Delta=1/k$. Since $\Delta(a^i)$ is a one-to-one function of the action frequency, then the the fact that asymptotically $\Delta=1/k$ a.s. implies that the action frequency a.s. converges. That the limit is the correct action frequency now follows form Lemma 3. *Q.E.D.*

Example 2, continued. It follows directly from Proposition 5 (part 2) that in the linear-absolute-cost model (and its asymmetric generalization mentioned earlier), the action frequency converges almost surely to the correct one. ■

Example 1, continued. Assume, without loss, that the true state is $\theta = 1$. In the linear-proportional-cost model with $k < 1$, part 1 of Proposition 5 implies that a herd occurs almost surely and there is a positive probability that the limit frequency of action 1 is one and a positive probability that it is zero, hence with positive probability is not correct.

When $k \geq 1$, it follows directly from Proposition 5 (part 3) that the limit frequency of action 1 is the correct frequency, $y^* = \frac{1+k}{2k}$. When $k = 1$, only action 1 is played in the long run. When $k > 1$, despite complete learning, both actions are played in the long-run. For every finite k, action 1 is played more often than action -1, but as k grows, this difference decreases and, in the limit as $k \to \infty$, there is an equal proportion of agents who choose both actions. Intuitively, as k increases (given $k \geq 1$) a smaller proportion of agents need to choose action 1 before a subsequent agent will switch to action -1 due to congestion. Interestingly, when $k \to 0$ the action frequency in the linear-proportional-cost model converges to the same level that is obtained, for arbitrary k, in the linear-absolute-cost model. ■

5.2 Choosing the superior action

It is natural to 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. This statistic can also be compared with its counterpart in the standard model without congestion.[24] In particular, we will see how a social planner in the standard model without congestion can use transfers to obtain the first-best outcome asymptotically by mimicking suitable congestion costs.

Under autarky, more than half of all agents choose the superior action. Because each agent in the model without congestion has more information than under autarky, the proportion of agents who choose the superior action exceeds that under autarky. The next proposition investigates the proportion of agents who choose the superior action under different forms of congestion costs.

Proposition 6.

1. *If total congestion can get large and there is constant incremental congestion, then, ex-ante (before the state is drawn), the proportion of agents who eventually choose the superior action is* $1/2$.
2. *If total congestion can get large, there is vanishing incremental congestion and* $\Delta(a^i)$ *can be written as one-to-one function of* $y(a^i)$*, then, ex-ante, the proportion of agents who*

[24] Comparing equilibrium utilities (even asymptotically) across the two settings is less compelling. If one views players' utility functions as just representing their preferences, then the comparison is not very meaningful since preferences differ. Furthermore, congestion costs affect a player's behavior entirely through the difference in congestion cost a player 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.

eventually choose the superior action is $1/2 + (1/2)[\Delta^{-1}(1/k) - \Delta^{-1}(-1/k)]$.

Proof. Part 1: From Part 2 of Proposition 5 and Part 2 of Lemma 3, the action frequency converges almost surely to $N/(1+N)$, regardless of the realized state. So, ex-ante, the proportion of agents who choose the superior action is $\frac{1}{2}\frac{1}{1+N} + \frac{1}{2}\frac{N}{1+N} = 1/2$.

Part 2: From Part 3 of Proposition 5 and Part 3 of Lemma 3, the action frequency converges almost surely to $\Delta^{-1}(1/k)$ when the realized state is 1, and to $\Delta^{-1}(-1/k)$ when the realized state is -1. So, ex-ante, the proportion of agents who choose the superior action is $(1/2)\Delta^{-1}(1/k) + (1/2)(1 - \Delta^{-1}(-1/k)) = 1/2 + (1/2)[\Delta^{-1}(1/k) - \Delta^{-1}(-1/k)]$. *Q.E.D.*

Although congestion costs may enable complete learning, the first part of Proposition 6 shows that congestion costs do not necessarily lead agents to choose the superior action more frequently than they do in the standard model without congestion, even in very natural environments. For instance, in the linear-absolute-cost model, half of all players choose the inferior action, more than would do so in the model without congestion. Indeed, in this case, fewer players choose the superior action than do so under autarky!

The second part of Proposition 6 implies that, in some environments, the proportion of agents who choose the superior action converges to 1. Consider, for instance, the linear-proportional-cost model when total congestion is large, $k \geq 1$. There, the proportion of agents choosing the superior action is $(1+k)/2k$ when the state is 1, and $(k+1)/2k$ when the state is -1, and therefore, the ex-ante proportion of agents choosing the superior action is $(1+k)/2k$. 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 goes to ∞.

An interesting corollary is that in the canonical observational learning model without congestion, a social planner could use a very simple transfer scheme to ensure that asymptotically, all agents will choose the superior action. This is achieved by effectively creating congestion costs that take the form of the linear proportional model with $k = 1$. For example, a social planner could require agent i to pay an amount $\tau_i(a^{i-1}) = y(a^{i-1})$ if he chooses action 1, and pay $\tau_i(a^{i-1}) = 1 - y(a^{i-1})$ if agent i chooses action -1. The transfer of agent i can be redistributed as a form of subsidy to subsequent agents (independent of their choices). Under this simple transfer scheme, regardless of the realization of the state, in the long-run the fraction of agents who take the superior action is one.

6 Applications

6.1 Pricing as congestion

One simple application of our results is to a problem 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 being a high-quality product by the state $\theta = 1$, whereas $\theta = -1$ represents product B being of high quality. Gross of price, the value of the high-quality product to any consumer is 1 while the low-quality product yields 0. Denote the decision to purchase product A by $a_i = 1$ and the decision to purchase product B as $a_i = -1$. Consumers make purchase decisions sequentially and each consumer observes the history of purchase decisions.

Assume that product pricing is as follows: after history a^i, the price of product A 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 Zemsky (1998), but with an arbitrary distribution of signals of bounded informativeness. The parameter range $k < 1$ can be interpreted as capturing 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 p chooses $a_{i+1} = 1$ (i.e. buy A) if and only if $r(p, q(a^i)) - kq(a^i) > 1 - r(p, q(a^i)) - k[1 - q(a^i)]$, or equivalently,

$$r(p, q(a^i)) > 1/2 + (k/2)(2q(a^i) - 1). \tag{7}$$

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 $\Delta(a^i) := 2q(a^i) - 1 = \frac{1-l(a^i)}{1+l(a^i)}$ so that the posterior belief threshold of our general model (Lemma 1), $r^*(\cdot) = 1/2 + k\Delta(\cdot)/2$, coincides with threshold implied by (7).[25] Note that $\Delta(a^i)$, as just defined, satisfies Assumption 1 and Assumption 2. We now consider two cases.

Case A. Suppose $k = 1$. Since private beliefs are informative (i.e. $\overline{b} > 1/2 > \underline{b}$), we have that for any a^i, $r(\overline{b}, l(a^i)) > q(a^i)$. Consequently, $\overline{\Delta}(\cdot; 1)$ has the property that for any $l > 0$, there exists $\overline{x}(l) > 0$ such that if $l(a^i) = l$ then

$$\Delta(a^i) \leq \overline{\Delta}(l; 1) - \overline{x}(l).^{26} \tag{8}$$

Similarly, for any $l > 0$, there is exists $\underline{x}(l) > 0$ such that if $l(a^i) = l$ then $\Delta(a^i) \geq \underline{\Delta}(l; 1) - \underline{x}(l)$. It follows that actions are always informative on the path of play, hence $\lim \Delta(a^i, 1, 1, ...) = 1$ and total congestion can get large. Now suppose that complete learning fails. Then, there is some $x > 0$ that is in $\mathrm{Supp}[l^k_\infty]$. Following the argument used in the proof of Theorem 2, for any $\varepsilon > 0$ there must be some a^i and some t such that (5) with $k = 1$ holds. But as $\varepsilon \to 0$, $\Delta^*(x, \varepsilon; 1) \to \overline{\Delta}(x; 1)$ (recall the proof of Theorem 1), which contradicts (8).

Case B. Suppose $k < 1$. Then, since $1/k > 1$ while $\Delta(\cdot) \in (-1, 1)$, the setting satis-

[25] Even though the $\Delta(a^i)$ function thus defined 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.

[26] In particular, one can choose $\overline{x}(l) = \overline{\Delta}(l; 1) - \Delta(a^i) = \frac{\overline{b}-(1-\overline{b})l}{\overline{b}+(1-\overline{b})l} - \frac{1-l}{1+l}$.

fies small total congestion. Hence, Corollary 1 implies that there is bounded learning, and the first part of Proposition 5 that there will be herd, which may be correct or incorrect. Indeed, since the setting has bounded total congestion (because $\Delta(a^i) = 2q(a^i) - 1$ implies that $\lim \Delta(a^i, 1, 1, ...) \leq 1 < 1/k$ and $\lim \Delta(a^i, -1, -1, ...) \geq -1 > -1/k$), there is bounded learning in the no-congestion limit.

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. In a setting with binary signals, there is an equilibrium in which the only people who play informatively are those who move first after a price change (plus the initial mover). Another variation on Avery and Zemsky (1998) limits prices to lid on a grid, e.g. pounds and pence. Because this restriction satisfies gaps, Theorem 3 implies bounded learning.

6.2 Queuing as congestion

We modify the linear-absolute-cost model of Example 2 to incorporate the idea that more recent predecessors exert larger effect on congestion costs than do more distant predecessors.

In the *queuing model with constant, unobservable service rate*, $c(a^1) = 0$ and for each $i \geq 2$,

$$c\left(a^i\right) = \sum_{j=1}^{i-1} \mathbb{1}_{\{a_j = a_i\}} \delta^{i-j-1},$$

for some $\delta \in [0, 1)$. In the extreme case where $\delta = 0$, every player cares only about her immediate predecessor's action. When $\delta > 0$, congestion depends more strongly upon recent than upon distant predecessors' choices. This cost function naturally generalizes a queueing model to allow for a constant service rate. For instance, player k may have observed whether each player $j < k$ entered restaurant $a = -1$ or restaurant $a = 1$ through the front door. However, she may not know whether j was served and exited through the restaurant's unobservable back door or remains. If she is risk neutral and believes that a $j < k - 1$ who remains in the restaurant exits every period with probability $1 - \delta$, and player $k - 1$ remains for sure, then she faces the congestion costs described above.

We observe that for any a^i,

$$\Delta(a^i, 1) = 1 + \delta\Delta(a^i) \text{ and } \Delta(a^i, -1) = \delta\Delta(a^i) - 1.$$

Furthermore, for any a^i, $\lim \Delta(a^i, 1, 1, \ldots) = \frac{1}{1-\delta}$, and $\lim \Delta(a^i, -1, -1, \ldots) = -\frac{1}{1-\delta}$. Hence, total congestion is bounded. Total congestion can get large when $k \geq 1 - \delta$, whereas it is small when $k < 1 - \delta$. For $\delta = 0$, it is easy to check that there is constant incremental congestion; hence, congestion has gaps. These properties, together with Theorem 1, imply the following:

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

We now consider the case 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 binary 0-1 random variable $S_j(t)$ to equal 1 with probability 1 when $t = j+1$, and for each $t \geq j+1$,

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

$S_j(t)$ is an indicator variable that describes whether Player j remains unserved through period t. Player i faces the congestion cost

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

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

Observe that for any $\delta > 0$ and any $k > 0$, congestion costs get arbitrarily large with probability one (e.g., when $S_j(\cdot) = 1$ 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 player must play informatively, but rather that it is always inevitable that some player in the future will do so.

7 Conclusion

We have investigated 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 deceases when more of his predecessors choose that option. This feature arises naturally in markets through changing prices, and in other environments where costs may stem from delays in service. Our analysis illustrates how different forms of congestion costs impact long-run learning and action frequency.

While we have focussed on settings where an agent finds an action more attractive when it has heretofore been rare, our approach can also be used to analyze situations where the direction of the externality is reversed. 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.[27]

[27] Under mild conditions one can also show that there is a.s. an information cascade. For example, the following

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.[28] More interestingly, with large total congestion, complete learning may fail, in contrast to the model without congestion. For example, in the linear absolute-cost model, when $k > 0$ is sufficiently small, complete learning fails even under unbounded private beliefs. On the other hand, there is still learning with high probability in the no-congestion limit so long as incremental congestion is bounded, as in the linear absolute-cost model.

A key assumption in this paper is that only past actions affects the payoff of an agent. There are, of course, cases where future actions also matter. Consider the case of betting, where each bettor i bets on which of two locations contains a prize. In some fixed-odds systems used by bookmakers, each bettor receives odds that depend on how many prior bettors have chosen each location, 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 location 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.[29]

Finally, we wish to emphasize the role that full rationality plays in the results. To see this, consider the contrast with Eyster and Rabin (2010)'s simple model of "naïve herding", in which players misconstrue their predecessors as basing their behavior solely upon their own private signals, failing to recognize that their predecessors perform informational inference as well. Eyster and Rabin (2010) show that not only do naïve players herd in a stronger sense than rational ones, but they also do so in a much broader set of environments. Indeed, unlike with rational players, congestion costs generally do not enhance social learning with naive players. Consider, for example, the linear-absolute-costs model with binary signals and small k. Once two more private signals favor action $a \in \{-1, 1\}$ over action $-a$, a temporary herd forms on a, during which the naïve players interpret each a choice as an additional private signal in favor of a. As k is small, they develop strong confidence in a being superior before congestion induces switching. In fact, because naïve players expect their predecessors to base their behavior solely on congestion and their private binary signal, they expect predecessors to start playing $-a$ uninformatively long before people actually do. If we make the assumption that a naïve player surprised to see her predecessor choose a never interprets that as evidence in favor of $-a$, then the strong confidence in a being superior will not be weakened at all before congestion drives people to finally break the herd, and the congestion-induced break in the herd will be deemed uninformative. In the limit as $k \to 0$, therefore, the probability that naïve

condition will be sufficient for an information cascade: there exists $\varepsilon > 0$ such that for any $(a^i, 1, 1, ...)$ and any $j > i$, $\Delta(a^j) > \varepsilon$ (and analogously if there is a herd on -1), where $\Delta(a^j)$ is now the net benefit that individual j derives from taking action 1 rather than -1.

[28] If private beliefs are unbounded we have that $\overline{\Delta}(l; k) = 1/k$ and $\underline{\Delta}(l; k) = -1/k$, and hence an agents's action is always informative under small total congestion.

[29] Koessler, Noussair and Ziegelmeyer (2008) make some progress with characterizing equilibrium behavior for sequential parimutuel betting with a very small number of bettors.

players become fully confident in the wrong state equals the likelihood that rational players herd on the wrong action without congestion.

This once more clarifies the essential driving force of the rational-learning results in this paper: small incremental but large total congestion guarantees that aggregate behavior repeatedly returns to situations where marginal behavior is informative. Our result relies on the correct reading of that informational content. Departures from full rationality might not completely reverse this result; for example, if players made random errors in understanding the meaning of congestion-induced behavior, we conjecture the results would be much the same. But systematically misconstruing that informational content is likely to undermine these results. In the case of naivety, because players infer too little from those avoiding congestion, players may never unlearn the lessons from the original crowds. If players somehow exaggerated the meaning of avoiding congestion relative to living with congestion, beliefs might never converge.

References

Acemoglu, Daron, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar, "Bayesian Learning in Social Networks," 2011. forthcoming in the *Review of Economic Studies*.

Ali, S. Nageeb and Navin Kartik, "Herding with Collective Preferences," 2011. Forthcoming in *Economic Theory*.

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.

Dasgupta, Amil, "Social Learning with Payoff Complementarities," May 2000. London School of Economics.

Debo, G. Laurens, Christine Parlour, and Rajan Uday, "Inferring Quality from a Queue," *Managerial Science*, 2012.

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 Roider, "The Wisdom of the Minority," *International Journal of Industrial Organization*, 2007, *25*, 391–415.

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

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.

_ , _ , _ , and Richard McKelvey, "Self-Correcting Information Cascades," *Review of Economic Studies*, 2007, *74* (3), 733–762.

Herrera, Helios and Johannes Hörner, "A Necesssary and Sufficient Condition for Information Cascades," 2011. 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.

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.

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.