

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Cripps, Martin; Thomas, Caroline D.

Article Strategic experimentation in queues

Theoretical Economics

Provided in Cooperation with: The Econometric Society

Suggested Citation: Cripps, Martin; Thomas, Caroline D. (2019) : Strategic experimentation in queues, Theoretical Economics, ISSN 1555-7561, The Econometric Society, New Haven, CT, Vol. 14, Iss. 2, pp. 647-708, https://doi.org/10.3982/TE1796

This Version is available at: https://hdl.handle.net/10419/217056

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.



https://creativecommons.org/licenses/by-nc/4.0/

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU

Strategic experimentation in queues

MARTIN W. CRIPPS Department of Economics, University College London

CAROLINE D. THOMAS Department of Economics, University of Texas at Austin

We analyze the social and private learning at the symmetric equilibria of a queueing game with strategic experimentation. An infinite sequence of agents arrive at a server that processes them at an unknown rate. The number of agents served at each date is either a geometric random variable in the good state or zero in the bad state. The queue lengthens with each new arrival and shortens if the agents are served or choose to quit the queue. Agents can observe only the evolution of the queue after they arrive; they, therefore, solve a strategic experimentation problem when deciding how long to wait to learn about the probability of service. The agents, in addition, benefit from an informational externality by observing the length of the queue and the actions of other agents. They also incur a negative payoff externality, as those at the front of the queue delay the service of those at the back. We solve for the long-run equilibrium behavior of this queue and show there are typically mass exits from the queue, even if the server is in the good state. KEYWORDS. Experimentation, bandit problems, social learning, herding, queues. JEL CLASSIFICATION. C72, C73.

1. INTRODUCTION

This paper considers a model that combines individual experimentation, observational learning, and payoff externalities. Such a combination arises in many economic and social contexts. Consider firms engaged in research and development (R&D) projects in related areas. If one firm has a success, this is good news for other firms, since it indicates that the entire area of research is worthwhile. However, the greater is the number of firms that are competing in the area, the less lucrative is the value of any patent that the firm secures. Similar concerns arise in other contexts, such as firms drilling for oil in the same geographical area or lenders to venture capitalists in a nascent industry.

The model we analyze has a countable infinity of individuals who arrive sequentially and join a queue for service. The queue grows at each new arrival and shrinks if service occurs or if an individual decides to stop waiting and leaves. Individuals arrive uncertain

Caroline D. Thomas: caroline.thomas@austin.utexas.edu

Martin W. Cripps: m.cripps@ucl.ac.uk

Cripps thanks the Cowles Foundation for its hospitality. Thomas gratefully acknowledges support from Deutsche Bank through IAS Princeton. We are grateful to V. Bhaskar, Max Stinchcombe, and Tom Wiseman for their comments. We thank four remarkably constructive referees for their comments and suggestions.

^{© 2019} The Authors. Licensed under the Creative Commons Attribution-NonCommercial License 4.0. Available at http://econtheory.org. https://doi.org/10.3982/TE1796

about the state of the world, which determines whether service occurs. Service occurs with positive probability only in the good state of the world. An individual has no direct information about events taking place before her arrival, but learns subsequently by observing the server activity and the behavior of other agents in the queue.

This is a situation faced by most of us as we approach counters for service or as we wait for taxis in unfamiliar places. Are short taxi queues a good sign because they indicate a high service rate, or a bad sign because informed individuals know that it is better to walk or take public transport?

As she waits in line without observing service, an individual revises downward the likelihood she attributes to service ever occurring. This is the usual private learning that occurs in strategic experimentation models. Additionally, the behavior of other agents in this game is itself a source of information. Social learning takes two different forms. Once in the queue, an individual learns from the behavior (leave or keep queueing) of those ahead of her in the queue. For instance, observing an agent ahead of her leave the queue is bad news about the state of the world. A more subtle aspect of social learning arises from the inference drawn by an arriving individual when she observes the queue length. Given a strategy profile, the server state determines an invariant distribution over queue lengths. Thus, the queue length observed on arrival is also informative about the server state.

While the presence of other individuals—those in the queue who are ahead of her is beneficial for an individual since they generate positive information externalities, it is detrimental in terms of payoffs, as the agent must wait for them to be served before she is.¹ Thus our model combines observational learning and payoff externalities in a novel manner.

Our main results are as follows. We study equilibria in a class of strategies that combine herding and individual experimentation in a natural way.² The first agent in line engages in optimal individual experimentation, choosing the amount of time she waits for service before quitting the queue and taking the outside option. Later arrivals in the queue copy the decisions of the first in line and quit whenever she quits. For each agent, the decision whether to join the queue or balk upon arrival is determined by a threshold queue length, and agents agree to join the queue only up to that threshold. We provide sufficient conditions for the existence of a symmetric equilibrium in such strategies.

For later arrivals, quitting the queue when the first in line does is optimal in equilibrium because the agents in the queue typically have nested information sets. Those at the front of the line have spent more time in the system and know strictly more than those behind them. This gives a novel equilibrium prediction: there is mass quitting of the queue, even when the server is good.

We show that equilibria can take two qualitatively different forms. When agents are sufficiently patient and/or good queues have sufficiently fast service, agents are willing to join the queue even when it is very long. In equilibrium, the persistence of the first in line reveals that she is sure that the state is good, and long queues can thus *perfectly*

¹We study first-come-first-served (FCFS) queues.

²We do not assume that deviations must be within that class.

reveal the good state of the world to new arrivals. At a good server, the system will typically alternate between successions of uninformed queues in which the agents renege en masse once the uninformed first in line becomes too pessimistic about the server state, and lengthy spells in which all individuals in the queue infer that the server is good and are willing to let the queue grow long. These perfectly informed spells end if the entire queue is served. At a bad server, we would only observe the cycle of uninformed queues.

In contrast, agents are unwilling to let the queues grow long when they are less patient and/or good queues are slow. In this case, there is no queue length that perfectly reveals the state to new arrivals, even if the first in line knows that the server is good. We then say that the equilibrium exhibits *imperfect revelation*. The equilibrium behavior of the queue at a good and at a bad server are then very similar. At a bad server, the queue cyclically repeats the following behavior. It grows to a limited length and stagnates there for a while before collapsing in a mass exit. At a good server, the queue would stagnate longer in that second stage if the individuals in the queue know that the state is good, but no further individuals are willing to join the queue.

This interplay of private and social learning means that, at our equilibria, relatively quiet periods of private learning alternate with bursts of activity at which social learning triggers herds. As a result, the queue lengths in our equilibria are more variable than in the cases where agents do not behave strategically or if the behavior of other agents in the queue were not observable. Hence, if designers care about the variability of the queue length, they may wish to limit the visibility of the queue.

Finally, it is important to note that at our equilibria, learning never ends. Every time the queue clears—either because the entire queue is served or because of mass exit—the next individual to arrive is unable to free-ride on the actions of better-informed individuals further up the line. The "social memory" is reset and individuals have to re-learn what past generations may already have learnt. Consequently if queues tend to empty out frequently, then on average there is poorer retention of past information. If queues tend to fill up, then social memory is improved and it is less necessary for new arrivals to duplicate past generations' learning.

The paper is organized as follows. In Section 2, we set up our queuing model. In Section 3, we introduce two concepts in the context of two auxiliary individual optimization problems. These concepts are then used to analyze the game.

Section 4 provides the main results of this paper. Equilibria with perfect revelation exist when agents are sufficiently patient. We derive necessary and sufficient conditions for the existence of a particular equilibrium with imperfect revelation, and show that, more generally, an equilibrium with imperfect revelation cannot exist if agents are too patient. These results are summarized in Proposition 1 in Section 4.3.

In Section 7, we discuss sufficient conditions that exclude the existence of equilibria in more general herding strategies. In Section 8, we show that faster servers encourage individual learning but are worse for social memory. Section 9 concludes and provides directions for further research.

1.1 Related literature

Queues are a pervasive feature of modern life and individuals join queues even when there is uncertainty about the service rate. When joining such a queue, an individual engages in optimal experimentation: while standing in line, she can learn about the arrival rate of service and, if this is poor, quit the queue and take her outside option. This problem also arises in many noneconomic situations³ (queueing for service in computer and communication networks, pipeline scheduling).

In addition to this private experimentation, individuals in the queue can see the behavior of others and this is informative about the service rate. The presence of other agents generates information externalities. Moreover, it is well known that in queues operating under a first-come–first-served (FCFS) regime, an individual who decides to join the queue imposes a negative payoff externality on those behind her (see Naor (1969) and Hassin (1985)).

This paper attempts to consider both types of externalities simultaneously. While combining these externalities leads to many analytical difficulties in general, queues provide a tractable structure within which this problem can be tackled. Strategic experimentation with information externalities has been widely studied (Bolton and Harris (1999), Keller et al. (2005), Murto and Välimäki (2011)). More recently, Strulovici (2010) and Thomas (2018) consider games of experimentation with direct payoff externalities.

Our model departs from much of the literature on herding and social learning⁴ in that we assume that each agent observes only the actions of other agents after she joins the queue.⁵ The queue length she observes upon arrival is her only summary statistic of previous behavior. Models assuming that agents do not observe the entire history of previous actions include Banerjee and Fudenberg (2004) and Smith and Sorensen (2013), where each new arrival observes only a random sample from the set of past observations, or Çelen and Kariv (2004), where each new arrival observes only her predecessor's action.

Research closely related to ours considers herding and social learning in the context of queues. Debo et al. (2012) consider a model in which the length of a queue reveals agents' private information about the unknown quality of a product for sale and explore a firm's incentive to manipulate the service rate. In this model, the equilibrium joining strategy is not of the threshold type: an uninformed individual joins a queue only if its length is strictly below or strictly above a certain threshold. In the latter case, the queue length reveals that the product is of high quality. Thus, as in our model, long queues can be a fully informative signal of the good state of the world. In a similar setup, Debo et al. (2018) explore a firm's incentive to manipulate prices. They prove the existence of a pooling equilibrium in which, even though the high and low quality firms set the same price, the queue grows long only when the quality is high. Eyster et al. (2014) study herding when a sequence of agents have the choice between two actions and bear a congestion cost determined by how many agents have previously chosen the action.

³See Percus and Percus (1990) or Chaudhry and Gupta (1996) for examples.

⁴See Banerjee (1992), Bikhchandani et al. (1992), Smith and Sørensen (2000), Callander and Hörner (2009). See Chamley (2004) for a survey.

⁵We thank an anonymous referee for this observation.

In all these models, the underlying uncertainty is about the value of service. As a result, all learning is done prior to the individuals' decision whether to join the queue (in Debo et al. (2012) and Debo et al. (2018)) or which action to take (in Eyster et al. (2014)). Once an individual has made this decision, there is no further learning, public or private. As a consequence, there is no experimentation and no reneging.

Reneging is an important feature of "real life" queuing decisions, and it occurs frequently in the equilibrium of our model—even when the state of the server is good. In the operations research literature on strategic behavior in queues (see Hassin and Haviv (2003) for a summary), a few articles provide an explanation for reneging. The model closest to ours is that of Mandelbaum and Shimkin (2000), where, with some probability, an arriving customer is diverted from an M/M/s queue and directed to a faulty server where she will never get served. Just as in our model, agents are uncertain about the service speed. The main difference is that in their model, the queue is not observable. An agent in the queue knows only that she has not yet been served, but she does not know whether she has been diverted. She also does not observe whether others have been served or whether others have left the queue. Because of this, a quitting agent would never trigger a herd. Nevertheless, the fact alone that she spends time in the queue without being served causes an agent to revise upward the probability she attributes to having been diverted, and makes her more inclined to renege. In equilibrium an individual might renege even if she has not been diverted.

An altogether different way to motivate reneging from an unobservable FCFS queue is to assume that the conditions in the system deteriorate exogenously. In Hassin and Haviv (1995), the value of being served exogenously drops to zero after some time T. In Haviv and Ritov (2001), individuals' waiting costs are increasing and convex with time. Shimkin and Mandelbaum (2004) allow for more general nonlinearities in the waiting cost function.

When the queue is observable, reneging from a FCFS M/M/s queue is harder to explain as equilibrium behavior. When other agents renege on the queue, an individual's prospects either remain the same or improve over time. In both cases, if it was worth joining the queue, then it is worth staying until service is completed. Assaf and Haviv (1990) and Altman and Shimkin (1998) show that reneging and balking may happen at egalitarian processor sharing systems (where the service capacity is split evenly among all agents present in the queue), because conditions in the system might deteriorate endogenously due to a slowdown in the service rate as more and more individuals join the queue. To our knowledge, our model is the first in which reneging is caused endogenously at an observable FCFS queue, namely through private and public learning about the service rate.⁶

When the entire queue is served and clears or if all agents in the queue leave en masse, new arrivals find that there are no agents from which to learn. These events happen with positive probability so that in this game learning never stops. The frequency of the events resetting social learning measures the persistence of social memory. Similar issues, although in a different context, are discussed in Herrera and Hörner (2013).

⁶We thank two anonymous referees for this observation.

Finally, in our equilibria, information can aggregate "in waves": in between informational cascades and ensuing herds, there will be periods of relative inactivity during which learning occurs gradually. Our model shares this feature with Bulow and Klemperer (1994), Toxvaerd (2008), and Murto and Välimäki (2011).

2. The model

Time is discrete, doubly infinite, and indexed by $\tau \in \mathbb{Z}$. At each date τ , one new agent arrives at the queue.⁷ The state of the server of our queue is either *good* or *bad*. The server is selected by nature once and for all at the outset of the game: nature selects the good server with probability $\mu \in (0, 1)$. A bad server never produces service capacity,⁸ and $g_{\tau} = 0$ for all τ , where $g_{\tau} \in \mathbb{N}^0 = \{0, 1, ...\}$ denotes the service capacity produced by the server at date τ . Only a good server produces service. In this state g_{τ} is an independent and identically distributed (i.i.d.) geometric⁹ random variable with commonly known parameter $\alpha \in (0, 1)$: $\Pr(g_{\tau} = i) = (1 - \alpha)\alpha^i$ for $i \in \mathbb{N}^0$.

Let $n_{\tau} \in \mathbb{N}^0$ be the number of agents in the queue at the *beginning* of period τ . The queue discipline is first-come–first-served (FCFS), that is, agents are served in the order of the queue. At each date τ , we distinguish three consecutive stages: service (*S*), exit (*E*), and arrival (*A*). The *S*, *E*, and *A* stages proceed as follows.

Service. If $n_{\tau} > g_{\tau}$, then the agents at the first g_{τ} positions in line are served at the service stage of date τ and disappear from the queue. Each remaining agent observes this and advances by g_{τ} positions. If $n_{\tau} \le g_{\tau}$, the entire queue is served and the excess service capacity, $g_{\tau} - n_{\tau}$, disappears. (It cannot be stored for use in subsequent periods.)

Exit. This is the only stage of date τ at which an agent still in the queue after the service stage may leave the queue and take the outside option. (This is called *reneging* on the decision to queue.) Any exit is observed by all agents who are still in the queue. Multiple rounds of exit are allowed at this stage.

Arrival. At this final stage, one new agent arrives and observes the number of agents remaining in the queue after the exit stage of date τ . This new agent can then choose either to join the queue at the last position or to take the outside option immediately. (Not joining the queue upon arrival is called *balking.*) The agent's decision is observed by all agents in the queue, although this information turns out not to matter. Once the arrival stage is concluded, the game moves to the next time period, $\tau + 1$.¹⁰

⁹This discrete-time geometric distribution service model for queues is widely used to model computer communication systems: see, for example, Chaudhry and Gupta (1996).

¹⁰The ordering of these stages appears to be the most natural. However, there are other possible modelling choices that could be made: If the *S* stage came after the *E* stage, then it would not be possible for

⁷Our timing choice allows us to clearly disentangle the two aspects of social learning from one another and from private learning. The assumption that arrival is deterministic simplifies our analysis of the agents' belief updating, as it ensures the existence of closed-form expressions for the state-dependent stationary distributions of queue lengths described in Section 5.1. See Wolff (1982) for properties of queues in continuous time.

⁸In an alternative model where, in the bad state, service does occur, but at a slower rate than in the good state, the steady-state behavior of queue lengths is more complicated. Since in such a model, observing service does not conclusively reveal that the server is good, an agent might eventually renege on the queue, even if she has previously observed service. This means that mass reneging can occur for many possible queue lengths.

While agents wait in line they receive a flow payoff of zero. An agent who is served obtains an instantaneous payoff of w > 1. Any agent who exits, either initially or after waiting for some time (balks or reneges), receives an instantaneous payoff normalized to 1. Each agent discounts each unit of calendar time by the common discount factor $\delta \in (0, 1)$.

The agent who arrives at date τ is uncertain about the state of the server. Each agent attaches prior probability μ to the server being in the good state. Agents in the queue at the start of period τ observe g_{τ} and the exit decisions of all agents at the exit stage. They must decide when, if ever, to renege on the queue and irrevocably take the payoff of 1. We analyze the steady state of this game.

In this model, agents must wait in line for at least one period before they have the opportunity to be served. The average rate at which agents are served in the good state is $\alpha/(1-\alpha)$. We define the parameter ϕ to be the inverse of the average service rate:

$$\phi := \frac{1 - \alpha}{\alpha}.\tag{1}$$

Much of the analysis below is done using ϕ rather than α . At a good server, if $\alpha > 1/2$ (or $\phi < 1$), the average service rate is greater than the arrival rate, and queues tend to empty. If $\alpha < 1/2$ ($\phi > 1$), queues tend to grow.

3. Two auxiliary optimization problems

We begin by describing the solutions to two individual optimization problems for an agent at the *n*th position in the line. In both of these optimizations, the behavior of all agents ahead of the *n*th in line is fixed: it is assumed that they remain in the queue until they are served.

The first optimization problem answers the following question: given that agents ahead of her never renege and she knows that the server is good, should an agent arriving at the *n*th position join the queue or balk and take the outside option? This individual problem is solved in Naor (1969). In our framework it determines \mathcal{M} , the maximum rational queue length at a server known to be good.

The second optimization is one of experimentation in the absence of social learning. It answers the following question: assuming that agents ahead of her never renege, how many periods should an agent at the *n*th position, holding the belief μ_n^0 on the good server, wait without observing service before reneging on the queue and taking the outside option? We interpret the solution to this problem, $\mathcal{N}(n, \mu_n^0)$, as the *n*th in line's willingness to experiment, based only on her private observation of the server (in)activity.

The variables \mathcal{M} and $\mathcal{N}(n, \mu_n^0)$ are used in later sections, when studying an agent's problem in our game of incomplete information.

those in the queue to exit immediately after bad news (no service) was observed; instead, they would need to wait one period before they could exit. This would reduce the equilibrium amount experimentation. If the A stage preceded the S stage, then new arrivals would immediately be able to observe the service behavior of the queue. This would make balking less likely and equilibrium queues would typically be longer. Finally, if the A stage preceded the E stage, then joining the queue would become more attractive as new arrivals could get to observe the exit behavior of others in the queue; again balking would be less common.

3.1 The maximal rational queue length at a good server: \mathcal{M}

Let V_n denote the expected payoff of an agent who knows that the server is in the good state, is *n*th in line following the current period's service stage, and assumes that all agents ahead of her remain in the queue until served. It satisfies the recursion $V_n = (1 - \alpha)\delta \sum_{t=0}^{n-1} \alpha^t V_{n-t} + \alpha^n \delta w$, with $V_1 = (1 - \alpha)\delta V_1 + \delta \alpha w$. Solving iteratively,

$$V_n = \psi^n \delta w$$
, where $\psi := \frac{\alpha}{1 - \delta(1 - \alpha)}$. (2)

Each additional agent ahead of her in the queue discounts the *n*th agent's payoff by the factor $\psi < 1$.¹¹ Hence, the parameter ψ captures the *congestion cost* imposed by an agent on those behind her. This congestion cost is mitigated as the service rate increases and it entirely disappears as α , and so ψ , approach unity. In contrast when service is slow, the congestion costs become extreme.

The threshold \mathcal{M} is the last position at which an agent accepts to join the queue at a server known to be good. It is the integer that satisfies

$$V_{\mathcal{M}+1} < 1 \le V_{\mathcal{M}}.\tag{3}$$

It depends¹² on the parameters α , δ , and w, but not the prior. The threshold \mathcal{M} grows without bound as congestion costs vanish, that is, as the agents become more patient $(\delta \rightarrow 1)$ or as the service rate increases $(\alpha \rightarrow 1)$. These arguments are summarized in the following lemma, where $\lfloor x \rfloor$ denotes the integer part of x.

LEMMA 1 (Naor (1969)). The last position at which an agent agrees to join the queue at a server known to be good is given by

$$\mathcal{M} = \left\lfloor \frac{-\ln(\delta w)}{\ln \psi} \right\rfloor. \tag{4}$$

(a) We have $\mathcal{M} \ge 1$ if and only if $V_1 \ge 1$. (b) For all $\alpha \in (0, 1)$, $\lim_{\delta \to 1} \mathcal{M} = +\infty$. For all $\delta \in [1/w, 1)$, $\lim_{\alpha \to 1} \mathcal{M} = +\infty$.

PROOF. Part (a) is immediate from (3). Part (b) follows from the definition of ψ .

3.2 The nth in line's private experimentation

We now turn to the *n*th in line's private learning, or *experimentation*, when she is uncertain about the server state and holds the belief $\mu_n^0 > 0$ that the server is good.¹³ In this optimization, we maintain the assumption that all agents ahead of her never renege on the queue. As a consequence, the *n*th in line cannot learn anything from the actions

¹¹Observe that $\psi^k = \mathbb{E}(\delta^{\tilde{\tau}_k})$, where the random variable $\tilde{\tau}_k$ is the arrival time of the *k*th service event, for $k \in \mathbb{N}$.

¹²For notational convenience, we do not make the dependence on these parameters explicit.

¹³In Section 5.2 we describe how this belief is obtained at an equilibrium.

of those ahead of her: there are no informational externalities. Hence, after *m* periods without service, she reduces her belief that the server is good to μ_n^m , where

$$\mu_n^m = \frac{\mu_n^0 (1 - \alpha)^m}{1 - \mu_n^0 + \mu_n^0 (1 - \alpha)^m}, \quad \text{so that} \quad \frac{\mu_n^m}{1 - \mu_n^m} = (1 - \alpha)^m \frac{\mu_n^0}{1 - \mu_n^0}.$$
 (5)

The first successful service event resolves the uncertainty and she learns that the server is good (her belief jumps to unity), even though she herself may not immediately be served. This generalizes the usual bandit problem to one where arrival of good news does not immediately generate a reward. Here the reward (of service) for the agent arrives at some random time after the arrival of the good news.

We want to determine the length of time that an uninformed *n*th in line will optimally wait to learn the server state, given that those ahead of her never renege on the queue. Our first step is to evaluate W_n , the *n*th in line's expected payoff if the server is revealed to be good (that is, at least one agent is served) at the current service stage. For n = 1, this payoff is *w* and for n > 1, W_n is proportional to the value (V_{n-1}) of being n - 1th in line at a good server:

$$W_n = (1 - \alpha) \left(V_{n-1} + \alpha V_{n-2} + \dots + \alpha^{n-2} V_1 \right) + \alpha^{n-1} w = \delta^{-1} V_{n-1} = \psi^{n-1} w.$$
(6)

We can now describe the payoff $U_n(m, \mu_n^0)$ of an agent who joins the queue as the *n*th in line, has prior $\mu_n^0 > 0$, and adopts the following strategy. Wait *m* periods for a service event and if one occurs during these *m* periods, never leave the queue, but if no service is observed, then renege after the *m* periods of server inactivity. The details of $U_n(m, \mu_n^0)$ can be explained as follows. First, the agent expects to observe no service over *m* periods with probability $1 - \mu_n^0 + \mu_n^0(1 - \alpha)^m$. Second, she attaches probability μ_n^0 to the server being good and, in that case, probability $\alpha(1 - \alpha)^{s-1}$ to service first occurring in the *s*th period, for each s = 1, ..., m. She then receives the payoff W_n discounted by δ^s :

$$U_n(m,\mu_n^0) := \left(1 - \mu_n^0 + \mu_n^0 (1 - \alpha)^m\right) \delta^m + W_n \mu_n^0 \sum_{s=1}^m \delta^s \alpha (1 - \alpha)^{s-1}$$
(7)

$$= (1 - \mu_n^0)\delta^m + \mu_n^0 V_n - \mu_n^0 \delta^m (1 - \alpha)^m (V_n - 1).$$
(8)

The three terms on the right of (8) represent the agent's payoff from taking her outside option when the state is bad, her payoff from being served with certainty when the state is good, and a correction to this second term that allows for the possibility that she may be unlucky in the good state and not observe service in the *m* periods she waits.

In the absence of social learning, the agent who is *n*th in line solves the problem $\max_{m \in \mathbb{N}^0} U_n(m, \mu_n^0)$. Her optimal behavior can be described in terms of a cutoff posterior $\underline{\mu}_n$ such that it is optimal for the *n*th in line to experiment as long as $\mu_n^m > \underline{\mu}_n$ and to renege otherwise. Alternatively, it can be described in terms of $\mathcal{N}(n, \mu_n^0) := \arg \max_{m \in \mathbb{N}^0} U_n(m, \mu_n^0)$ as the number of service failures she observes before reneging. The lemma below describes both. Its proof is given in Appendix A.1.

LEMMA 2. For all $n \in \mathbb{N}$ and $\mu_n^0 \in (0, 1)$, there exists a solution, m^* , to the problem $\max_{m \in \mathbb{N}^0} U_n(m, \mu_n^0)$. The solution m^* is unique for almost every $\mu_n^0 \in (0, 1)$ and satisfies

$$m^{*} = \mathcal{N}(n, \mu_{n}^{0}) := \left[\left(\ln(1-\alpha) \right)^{-1} \ln\left(\frac{1-\mu_{n}^{0}}{\mu_{n}^{0}} \frac{\psi(1-\delta)}{\alpha(\psi^{n}\delta w - 1)} \right) \right]_{+}, \tag{9}$$

where $\lceil x \rceil_+$ denotes the smallest nonnegative integer greater than or equal to x. At this solution, the agent chooses to renege when her posterior hits the cutoff

$$\underline{\mu}_n = \frac{1-\delta}{\delta\alpha(\psi^{n-1}w-1)}.$$
(10)

For nongeneric beliefs $\mu_n^0 \in (0, 1)$ such that $U_n(m^*, \mu_n^0) = U_n(m^* + 1, \mu_n^0)$, it is optimal to experiment for m^* periods, or for $m^* + 1$ periods, or to randomize between the two.

4. Equilibrium of the queuing game

The next sections are organized as follows. First we describe the strategy σ^* (Definition 1) that will be played by every agent at an equilibrium of this game. Our equilibrium concept is the symmetric steady-state Bayesian equilibrium defined in Section 4.2. In equilibrium, the strategy σ^* induces two stationary distributions of queue lengths: one for each state of the server, good and bad. These are described in Section 5.1. When an agent arrives at the queue and observes its current length, she uses these distributions to form her posterior belief about the server state, as described in Section 5.2. If there exist queue lengths that perfectly reveal that the server is good, we say that the equilibrium strategy exhibits *perfect revelation*. If every queue length is only imperfectly revealing of the server state, we say that the equilibrium strategy exhibits *imperfect revelation*. Section 6.1 deals with equilibria with perfect revelation and Section 6.2 deals with those with imperfect revelation. Our results are summarized in Proposition 1 in Section 4.3.

4.1 Strategy

At each period, an agent's strategy maps her information into a binary choice, i.e., whether to be in the queue or to take the outside option.¹⁴ We consider the strategy $\sigma^*(q, N, M)$ defined below, where the probability $q \in (0, 1]$ and the nonnegative integers N and M are parameters of σ^* .

DEFINITION 1. The strategy $\sigma^*(q, N, M)$.

• Upon arriving at the queue, an agent joins the queue if and only if she is at most *M*th in line.

¹⁴A formal definition of a strategy is given in Appendix B.1 and a formal definition of the strategy $\sigma^*(q, N, M)$ is given in Appendix B.1.1, both in the Supplemental Material, available in the supplementary file on the journal website, http://econtheory.org/supp/1796/supplement.pdf.

- Once in the queue, if she observes service, she never reneges.
- Conditional on not observing service, two scenarios arise:
 - If she joined the queue at the first position, then she does not renege for the first N 1 periods. With probability $q \in (0, 1]$, she reneges at the exit stage of the *N*th period; with probability 1 q she reneges at the exit stage of the (N + 1)th period.
 - If there were agents ahead of her in the queue when she joined, then she reneges on the queue if and only if the first in line does, and in the same period as the first in line.

Under a strategy profile where every agent plays according to σ^* , no agent other than the first in line autonomously reneges on the queue. An agent joining the queue at the first position in line reneges after observing N (for q = 1) service failures. We say that the first in line *experiments* or is *uninformed* if she has never observed service and is therefore still uncertain about the server state. Since exactly one new agent arrives every period, the queue at a bad server can never be longer than N (or N + 1 for q < 1). The cases $N \le M$ and N > M are qualitatively different for q = 1 (for q < 1, the relevant cases are N < M and $N \ge M$).¹⁵

For $N \leq M$, agents continue to join the queue as long as the first in line experiments. At a bad server, the queue therefore grows to length N before the first in line reneges and the next period's arrival joins the queue at the first position. An agent joining a queue no longer than N does not know whether the agent currently first in line initially joined the queue at a later position and moved up to first position when service occurred or whether she joined the queue at the first position and has been waiting for service ever since. In other words, an agent joining a queue no longer than N does not know whether the first position and has been waiting for service ever since. In other words, an agent joining a queue no longer than N does not know whether the first in line has already observed service or is still experimenting.

This uncertainty is resolved once the queue reaches length N. If the first in line reneges after N periods without service, all those behind her infer that she had not yet observed service and that she was an uninformed first in line. If the first in line does not renege after N periods, all those behind her infer that she has previously observed service and that she is informed that the server is good. Those behind her are now certain that the server is in the good state, and will remain in the queue until served.

Similarly, an agent arriving at a queue in *n*th position, for $N < n \le M + 1$, can infer that the server is in the good state. In other words, the length of the queue itself is sufficient to reveal the first in line's information to agents arriving at positions n > N. We therefore say that the strategy σ^* exhibits perfect revelation when $N \le M$.

For M < N the queue never exceeds length M, even as the first in line continues to experiment. If the first in line does not renege after N unsuccessful service events, the agents queuing behind her learn that the server is in the good state. But even in that event the queue never grows longer than M. So while the position, n = 1, ..., M + 1, at which an agent arrives at the queue remains informative about the server state, there exists no n that conclusively reveals the server state. We say that the strategy σ^* exhibits imperfect revelation when M < N.

¹⁵Our discussions in the remainder of this paper refer to the case q = 1 unless otherwise specified.

4.2 Equilibrium

Consider the strategy σ^* defined in Definition 1 and consider the symmetric strategy profile where every agent plays according to the strategy σ^* .¹⁶ Under this profile the queue length is bounded by M. Then, for every possible server state, the strategy profile σ^* induces a finite Markov chain on the queue lengths. Consequently, σ^* determines a unique pair of state-dependent invariant distributions over queue lengths. The strategy σ^* , combined with the two server states, therefore defines a unique on-path¹⁷ belief system, $\mu(\sigma^*)$, where agents' beliefs regarding the underlying state of the server are defined at every information set reached with positive probability under σ^* , and are derived using Bayes' rule from the invariant distributions induced by σ^* .

We define a *symmetric steady-state Bayesian equilibrium* (henceforth equilibrium) to be a strategy σ^* played by every agent and an on-path belief system $\mu(\sigma^*)$ such that the following statements hold:

- E.1. Every belief in $\mu(\sigma^*)$ is derived using Bayes' rule from the unique pair of invariant distributions induced by σ^* .
- E.2. At every information set arising with positive probability on the path of play under σ^* , any action that is played with positive probability under σ^* maximizes the agent's payoff, given her beliefs in $\mu(\sigma^*)$ and given that every other agent adheres to σ^* .

Observe that E.2. does not require that deviations must be within the class of strategies of Definition 1. We wish to show that there exists (q^*, N^*, M^*) such that $\sigma^*(q^*, N^*, M^*)$ constitutes an equilibrium.¹⁸ For every (q, N, M), the state-dependent stationary measures of queue lengths induced by $\sigma^*(q, N, M)$ are described in Section 5.1. Let $\bar{\mu}_n^t$ denote the belief about the server state formed in accordance with E.1. by an agent having arrived at the queue at the *n*th position and having observed *t* periods without service.¹⁹ The on-path inference under $\sigma^*(q, N, M)$ is described in Section 5.2.

¹⁹In our notation we suppress the dependence of the equilibrium posterior on σ^* but use upper bars to denote these values. We write $\bar{\mu}_1^0(q, N, M)$ when we wish to emphasize the dependence on (q, N, M).

¹⁶In a slight abuse of notation, we henceforth use σ^* to refer both to the individual strategy σ^* and to the symmetric strategy profile where every agent plays according to σ^* .

¹⁷Beliefs are determined at every information set reached with positive probability under the strategy profile σ^* and are not determined at information sets that are not reached under the strategy profile σ^* .

¹⁸ In our analysis, we can ignore beliefs off the equilibrium path for the following reason: a player's offpath behavior does not affect her payoff and does not affect the invariant distributions of queue lengths. Moreover, if $\sigma^*(q^*, N^*, M^*)$ is a symmetric steady-state Bayesian equilibrium, then we can find off-path beliefs to support $\sigma^*(q^*, N^*, M^*)$ as a perfect Bayesian equilibrium profile. In our game a perfect Bayesian equilibrium is a strategy profile and a system of beliefs about the server state. The strategy profile must be sequentially rational at every information set given the agents' beliefs about the server state. The beliefs about the server state are defined by Bayes' rule at every on-path information set and may be chosen arbitrarily at any off-path information set. We can specify off-path beliefs as follows. If the first in line reneges at any queue length other than N^* (for $q^* = 1$), the belief of an agent behind her in the queue drops to $\tilde{\mu}$ (unless that agent had previously observed service), where $\tilde{\mu}$ is chosen to be sufficiently low that it is sequentially rational for the queue to clear. (For instance, $\tilde{\mu} := \mu_1^{N^*}$.) If an agent observes someone other than the first in line renege on the queue, even though the first in line does not, then her belief remains unaffected by this deviation.

We now describe the conditions that E.2. imposes on q^* , N^* , and M^* . We distinguish the cases with perfect revelation and imperfect revelation. For the sake of exposition, let us focus on pure strategies, where $q^* = 1$.

4.2.1 *Equilibrium with perfect revelation* By definition, the strategy $\sigma^*(1, N^*, M^*)$ exhibits perfect revelation if there exist queue lengths that perfectly reveal that the server is good. For $q^* = 1$, this requires $N^* \leq M^*$. For each $n = 1, ..., M^* + 1$, consider the condition E.2. applied to the agent arriving at the queue at the *n*th position in line, when all other agents adhere to the strategy $\sigma^*(1, N^*, M^*)$.

For $n \ge N^* + 1$, the agent learns that the server is good upon arriving at the queue. Given this belief, it must be optimal for her to join the queue and never renege if $n \le M^*$ and to balk if $n = M^* + 1$. Thus, the agent faces the problem described in Section 3.1, and by Lemma 1 she joins the queue (and never reneges) if and only if $n \le M$. Thus, E.2. requires that M^* equals \mathcal{M} , defined in (4).

For n = 1, the agent forms the belief $\bar{\mu}_1^0$ upon arriving at the queue. Given this belief, it must be optimal for her to experiment for N^* periods. The agent faces the problem described in Section 3.2, and by Lemma 2, N^* is optimal for her if and only if it equals $\mathcal{N}(1, \bar{\mu}_1^0)$ defined in (9).

For $n = 2, ..., N^*$, the agent forms the belief $\bar{\mu}_n^0$ upon arriving at the queue. Given this belief, it must be optimal for her to join the queue and herd on the first in line's actions. Reneging on the queue when the first in line reneges is clearly optimal: once the first in line reneges, every agent behind her adopts her posterior belief and faces at least as much congestion as her. Conversely, if the first in line does not renege after N^* periods, those behind her learn that the server is good, and by Lemma 1 it is optimal for them to stay in the queue until served.

What remains is the condition that no agent in the queue wants to autonomously renege *before* the queue reaches length N^* , the point at which an uninformed first in line reneges, thereby revealing that she is uninformed. The following tension arises. Information arrives "faster" for those later in line than for the first in line: the agent who joins the queue at the *n*th position needs to wait only $N^* - n + 1$ periods to obtain the first in line's information acquired over N^* periods of experimentation. However, those later in line also face more congestion, and the *n*th in line's expected payoff if the server were known to be good might be too low to make it worth her waiting for the first in line's information.

Let $U_n^*(\bar{\mu}_n^0)$ denote the payoff under the strategy $\sigma^*(1, N^*, M^*)$ for an agent joining the queue at the *n*th position, as a function of her belief $\bar{\mu}_n^0$ upon joining the queue. The equilibrium condition (A.iv) in the next lemma says that for an agent joining the queue at position $n = 2, ..., N^*$, the equilibrium payoff $U_n^*(\bar{\mu}_n^0)$ must exceed the payoff from reneging autonomously after $m \le N^* - n$ service failures, which is given by $U_n(m, \bar{\mu}_n^0)$, defined in (7).

The above results are summarized in the following lemma.

LEMMA A. The pure²⁰ strategy $\sigma^*(1, N^*, M^*)$ with perfect revelation satisfies the equilibrium condition E.2. if

- $(A.i) \quad N^* \leq M^*,$
- $(A.ii) \quad M^* = \mathcal{M},$
- (A.iii) $N^* = \mathcal{N}(1, \bar{\mu}_1^0(1, N^*, M^*)),$

(A.iv)
$$U_n^*(\bar{\mu}_n^0(1, N^*, M^*)) \ge U_n(m, \bar{\mu}_n^0(1, N^*, M^*))$$
 for all $n = 2, ..., N^*$ and for all $m = 0, 1, ..., N^* - n + 1$.

The remainder of this section is dedicated to deriving the expression for $U_n^*(\bar{\mu}_n^0)$ in (12). Let $\bar{\eta}_n$ denote the probability that an agent who joins the queue at the *n*th position attaches to the event that the first in line has not previously observed service, conditional on the server being good.²¹ This is therefore the probability that the *n*th in line attaches to the event that she will renege together with the first in line once the queue reaches length N^* , in which case her payoff is 1. If the first in line has previously observed service (probability $1 - \bar{\eta}_n$) so that she does not renege once the queue reaches length N^* , the *n*th in line learns that the server is good and her payoff is V_n , defined in (2). Consequently,

$$A_n = \bar{\eta}_n + (1 - \bar{\eta}_n) V_n \ge 1$$
(11)

is the *n*th in line's expectation upon joining the queue, conditional on the server being good, of the payoff she will obtain once the queue reaches length N^* .

We can now decompose $U_n^*(\bar{\mu}_n^0)$ as follows. The *n*th arrival expects the server to be bad with probability $1-\bar{\mu}_n^0$, in which case she reneges with the first in line after $N^* - n + 1$ service failures. She attaches probability $\bar{\mu}_n^0$ to the server being good, and in that case, probability $\alpha(1-\alpha)^{s-1}$ to service first occurring in the *s*th period, for each $s = 1, \ldots, N^*$ n + 1, yielding the payoff W_n discounted by δ^s . Finally, she attaches probability $\mu_n^0(1 - \alpha)^{N^* - n + 1}$ to the server being good, but producing $N^* - n + 1$ service failures. In this event, the queue reaches length N^* , and the first in line's behavior reveals whether she is informed or uninformed. The *n*th in line's resultant payoff is therefore A_n , defined above. In summary,

$$U_n^*(\bar{\mu}_n^0) = (1 - \bar{\mu}_n^0)\delta^{N^* - n + 1} + W_n\bar{\mu}_n^0\sum_{s=1}^{N^* - n + 1}\delta^s\alpha(1 - \alpha)^{s-1} + \bar{\mu}_n^0(1 - \alpha)^{N^* - n + 1}\delta^{N^* - n + 1}A_n.$$

The above simplifies to

$$U_n^*(\bar{\mu}_n^0) = (1 - \bar{\mu}_n^0)\delta^{N^* - n + 1} + \bar{\mu}_n^0 V_n - \bar{\mu}_n^0 (1 - \alpha)^{N^* - n + 1}\delta^{N^* - n + 1} (V_n - A_n).$$
(12)

²⁰If $q^* \in (0, 1)$, condition (A.i) becomes $N^* < M^*$ and (A.iii) becomes $N^* = \mathcal{N}(1, \bar{\mu}_1^0(q^*, N^*, M^*))$. Moreover, condition (A.iv) must hold for all $n = 2, ..., N^* + 1$ and for all $m = 0, 1, ..., N^* - n + 2$. This accounts for the fact that the *n*th in line obtains information from the first in line's behavior after $N^* - n + 1$ service failures, *and* after $N^* - n + 2$ service failures. Consequently, when $q^* \in (0, 1)$, $U_n(N^* - n + 1, \bar{\mu}_n^0)$ overestimates the *n*th in line's payoff from reneging after $N^* - n + 1$ service failures, for $n = 2, ..., N^* + 1$, and condition (A.iv) evaluated at $m = N^* - n + 1$ is sufficient, but not necessary for equilibrium.

 $^{^{21}}$ In equilibrium, this belief is formed in accordance with E.1., as described in (14) in Section 5.2.

4.2.2 *Equilibrium with imperfect revelation* By definition, the strategy $\sigma^*(1, N^*, M^*)$ exhibits imperfect revelation if every queue length is only imperfectly informative about the server state. For $q^* = 1$, this requires $N^* > M^*$.²² For each $n = 1, ..., M^* + 1$, consider the condition E.2. applied to the agent arriving at the queue at the *n*th position in line, when all other agents adhere to the strategy $\sigma^*(1, N^*, M^*)$.

For $n = M^*$, $M^* + 1$, the agent cannot learn, merely by observing the queue length, that the server is in the good state. Consequently, condition (A.ii) does not determine the equilibrium parameter M^* , as it did under perfect revelation. Instead, the conditions jointly pinning down M^* are given subsequently in (B.ii). They can be understood as follows.

First, an agent arriving at the queue at the M^* th position and forming the belief $\bar{\mu}_{M^*}^0$ must prefer joining the queue (thus adhering to the equilibrium strategy) to balking. Equivalently, her payoff from adhering to the equilibrium strategy, which is $U^*_{M^*}(\bar{\mu}_{M^*}^0)$ as defined in (12), must exceed 1, which is her payoff from balking.

Second, an agent arriving at the M^* + 1th position must find it optimal to balk (thus adhering to the equilibrium strategy). Equivalently, her payoff from balking must be no less than the payoff to her most profitable deviation from $\sigma^*(1, N^*, M^*)$. Some care is needed when describing her most profitable deviation, and the belief $\bar{\mu}_{M^*+1}^0$ she forms upon arrival. Indeed, observe that the agent arriving at the queue at the M^* + 1th position does not know whether she is the first, second, third,..., $N^* - M^*$ th agent to arrive at that position behind the current first in line. If she joins the queue, observing the first in line's behavior will be instructive regarding this point.

In particular, suppose that the agent arriving at the M^* + 1th position joins the queue and observes $N^* - M^*$ service failures. If the first in line does not renege after the $N^* - M^*$ th failure, the agent at the M^* + 1th position learns that she must have been the first arrival at the M^* + 1th position behind the current first in line. She also learns that the first in line has not reneged despite having observed a total of N^* service failures. This means that the first in line must have previously observed service. The agent at the M^* + 1th position therefore learns that the server is good. Consequently, it becomes optimal for her to remain in the queue until served.

Likewise, if at any point the agent observes service, it becomes optimal for her to remain in the queue until served. Finally, if at any point she observes the first in line renege, it is optimal for the agent at the M^* + 1th position to renege in the same period as the first in line.

Therefore, the most profitable deviation for the agent arriving at the M^* + 1th position must be one of the following strategies, parameterized by *m*.

• For $1 \le m \le N^* - M^*$, wait *m* periods. If a service event occurs during these *m* periods, remain in the queue until served. If the first in line reneges on the queue before any service event occurs, then renege in the same period as the first in line.

²²When $q^* = 1$, the strategy with $N^* = M^* > 1$ exhibits perfect revelation, since arriving at the $M^* + 1$ th position conclusively reveals that the server is good. When $q^* \in (0, 1)$, the strategy with $N^* = M^* > 1$ exhibits imperfect revelation, since arriving at the $M^* + 1$ th position does not conclusively reveal the server state. The strategy with $N^* = M^* = 1$ exhibits imperfect revelation for every $q^* \in (0, 1]$.

In all other cases (i.e., there are *m* failures and the first in line has not reneged), autonomously renege after *m* periods. Intuitively, the agent arriving at the $M^* + 1$ th position joins the queue but autonomously reneges, forgoing the chance to learn from an informed first in line that the server is good. We let $U_{M^*+1}(m, \bar{\mu}_{M^*+1}^0)$ denote the payoff to this strategy and give an expression for it in (A.56). The second condition in (B.ii) requires that this deviation be unprofitable for every $m = 1, \ldots, N^* - M^*$.

• Wait $m = M^* - N^*$ periods. If a service event occurs during these *m* periods, remain in the queue until served. If the first in line reneges on the queue before any service event occurs, then renege in the same period as the first in line. In all other cases (i.e., there are *m* failures and the first in line has not reneged), remain in the queue until served. Intuitively, the agent arriving at the $M^* + 1$ th position joins the queue and remains in line long enough that she is certain to obtain the first in line's information, either because the first in line reneges (this can happen after the $M^* + 1$ th in line has queued for $1, \ldots, M^* - N^*$ periods) or because the first in line does not renege after the $M^* + 1$ th in line has queued for $M^* - N^*$ periods, revealing that she is informed and that the server is good. Thus, the $M^* + 1$ th in line stays in the queue long enough to *herd* on the first in line. We let $U_{M^*+1}^h(\bar{\mu}_{M^*+1}^0)$ denote the payoff to this strategy, and give an expression for it in (A.59). The third condition in (B.ii) requires this deviation to be unprofitable.

Finally, for $n < M^*$, the agent's problem is the same as with perfect revelation. Conditions (B.iii) and (B.iv) below mirror conditions (A.iii) and (A.iv), respectively.

The above results are summarized in the following lemma.

LEMMA B. The pure²³ strategy $\sigma^*(1, N^*, M^*)$ with imperfect revelation satisfies the equilibrium condition E.2. if

$$(B.i) \quad N^* > M^*,$$

$$(B.ii) \quad \begin{cases} U_{M^*}^* (\bar{\mu}_{M^*}^0(1, N^*, M^*)) \ge 1, \\ \max_{m \le N^* - M^*} \mathcal{U}_{M^* + 1}(m, \bar{\mu}_{M^* + 1}^0(1, N^*, M^*)) \le 1, \\ \mathcal{U}_{M^* + 1}^h(\bar{\mu}_{M^* + 1}^0(1, N^*, M^*)) \le 1, \end{cases}$$

$$(B.iii) \quad N^* = \mathcal{N}(1, \bar{\mu}_1^0(1, N^*, M^*)),$$

(B.iv)
$$U_n^*(\bar{\mu}_n^0(1, N^*, M^*)) \ge U_n(m, \bar{\mu}_n^0(1, N^*, M^*))$$
 for all $n = 2, ..., M^*$ and
for all $m = 0, 1, ..., N^* - n + 1$.

²³If $q^* \in (0, 1)$, the following adjustments are necessary so as to account for the event that the first in line experiments for $N^* + 1$ periods. First, there are now $N^* - M^* + 1$ possible instances of the arrival at the $M^* + 1$ th position. Second, the agent arriving at the *n*th position, $n = 2, ..., M^*$, must now be willing to wait $N^* - n + 2$ periods to obtain the first in line's information. Third, an agent arriving at the $M^* + 1$ th position must prefer balking to waiting $N^* - M^*$ and to waiting $N^* - M^* + 1$ periods so as to herd on the first in line. In addition, strategies with $N^* = M^*$ exhibit imperfect revelation, and are therefore admissible. Hence, (B.i) becomes $N^* \ge M^*$. Finally, (B.iii) becomes $N^* = \mathcal{N}(1, \bar{\mu}_1^0(q^*, N^*, M^*))$.

4.3 Equilibrium existence

The main results of this paper are summarized in the next proposition. First we establish that an equilibrium with perfect revelation exists if agents are sufficiently patient. Second, if agents are too patient, an equilibrium with imperfect revelation cannot exist. Finally, equilibria with imperfect revelation exist for intermediate values of the discount factor.²⁴

- PROPOSITION 1. (i) Given any (α, μ) , there exists $\underline{\delta}(\alpha, \mu) < 1$ such that an equilibrium with perfect revelation exists for any $\delta \in (\underline{\delta}(\alpha, \mu), 1)$.
 - (ii) Given any (α, μ) , there exists $\overline{\delta}(\alpha, \mu) < 1$ such that an equilibrium with imperfect revelation cannot exist for any $\delta \in (\overline{\delta}(\alpha, \mu), 1)$.
 - (iii) Given any (α, μ) , there exists a pair of adjacent intervals $D_1^*(\alpha, \mu)$ and $D_2^*(\alpha, \mu)$ such that an equilibrium with imperfect revelation with $N^* = M^* = 1$ exists if and only if $\delta \in D_1^*(\alpha, \mu) \cup D_2^*(\alpha, \mu)$. If $\delta \in D_1^*(\alpha, \mu)$, then $q^* = 1$. If $\delta \in D_2^*(\alpha, \mu)$, then $q^* \in (0, 1)$.

The first part of this result is proved and discussed in Section 6.1; the remaining parts are proved and discussed in Section 6.2.

Figure 1 illustrates the equilibrium²⁵ values M^* and N^* as a function of the discount factor δ for three different values of α . Notice that there can be multiple equilibria (multiple values of N^* or M^*) for some (α, δ, μ) . For each α , equilibria with imperfect revelation exist only when the discount factor is small, whereas equilibria with perfect revelation exist only when the discount factor is large. Our figures illustrate that the transition is not sharp and that there are parameter values at which both types of equilibria exist.²⁶ When the discount factor goes to 1, M^* goes to infinity. The behavior of N^* as δ approaches 1 is discussed in Lemma 10 and depends on α . When $\alpha < 1/2$, N^* converges to 1; when $\alpha = 1/2$, N^* converges to a strictly positive constant; when $\alpha > 1/2$, N^* goes to infinity,²⁷ although at a slower rate than M^* .

In Section 7, we ask whether additional equilibria exist at which agents other than the first in line can trigger a herd. We show that for $\alpha \ge 1/2$ and δ sufficiently large, any given strategy with more than one herding leader cannot be an equilibrium.

²⁴If the discount factor is so low that $V_1 < 1$, then, at the unique equilibrium, all agents arriving at the queue immediately balk and take the outside option.

²⁵Figure 1 is based on a Mathematica simulation that focusses on pure strategy equilibria, except when $\alpha < 1/2$ and δ is so large that only mixed strategy equilibria with perfect revelation exist. These have $N^* = 1$ and $q^* < 1$: conditional on no service, the first in line is indifferent between reneging at the first or the second exit stage following her arrival. (See Lemma 10.)

²⁶For instance, when $(\alpha, \mu, \delta) = (0.7, 0.9, 0.6)$, there is an equilibrium with imperfect revelation that has $N^* = 4$ and $M^* = 3$, and an equilibrium with perfect revelation that has $N^* = 3$ and $M^* = 5$.

²⁷This effect is not visible in our figures due to the coarseness of the grid used on the interval of possible values of δ .



FIGURE 1. Equilibrium values M^* (open discs) and N^* (dots) as a function of δ for $\alpha = 0.7$, $\alpha = 0.5$, and $\alpha = 0.3$ (from top to bottom), and for $\mu = 0.9$. (The vertical lines emphasize values of δ at which there are multiple equilibria.)

5. ON-PATH BELIEFS

5.1 Stationary distributions of queue lengths

In equilibrium, the strategy $\sigma^*(q, N, M)$ determines the evolution of the queue length. There are two discrete-time Markov processes to consider: one that arises if the server is good and the other if it is bad. In this section we describe the stationary measure of each process. These will shape the inference of new arrivals as described in Section 5.2. We consider the stochastic process followed by the queue length *at the beginning of the arrival stage* of each date $\tau \in \mathbb{N}^0$. We say that the queue has length *n* at date τ if the agent arriving in the system at date τ arrives in the queue at the *n*th position—even if that agent then balks.

Let $(x_n(q, N, M))_{n=1}^{M+1} \in \Delta(M + 1)$ denote the stationary probabilities of arriving at the queue at the *n*th position under the strategy $\sigma^*(q, N, M)$, conditional on the server being in the bad state.²⁸ At a bad server, the queue follows an almost deterministic process. If $N \leq M$, the queue grows by one agent each period and then shrinks to length 1 in the period after reaching length N (if q = 1). Hence in the bad state, there is an ergodic probability $x_n = x_1$ of arriving *n*th in line (for $n \leq N$) and zero probability of arriving at the queue at any other position. If M < N, the queue grows to size M and then stays at that size for a further N - M periods before shrinking to unity. Thus there is a(nother) constant ergodic probability equal to x_1 of arriving at the M + 1st position (and balking).

At a good server, the queue length follows a more complex Markov chain, and the state of the process must be the position in the queue at which the latest agent arrives *and* whether or not the first in line knows that the server is in the good state. Allowing for q < 1, there are at most M + N + 2 states for this process: arrival at positions $1, 2, 3, \ldots, M + 1$ and the first in line knows the server is in the good state; arrival at positions $1, 2, 3, \ldots, M + 1$ and the first in line is uncertain. The process governing the queue length under the strategy $\sigma^*(q, N, M)$ at a good server has finite states and is irreducible. Therefore it must admit a unique stationary measure. We define $(y_n(q, N, M))_{n=1}^{M+1} \in \Delta(M + 1)$ to be the stationary probabilities of arriving at the *n*th position in line under the strategy $\sigma^*(q, N, M)$, conditional on the server being in the good state.

The exact expressions for $(x_n(q, N, M))_{n=1}^{M+1}$ and $(y_n(q, N, M))_{n=1}^{M+1}$ are given in Proposition A.1 in Appendix A.2. The stationary measure at a good server admits three qualitatively different forms depending on ϕ . These are described in the next subsection. Some readers may prefer to proceed directly to Section 5.2, which describes an agent's belief updating when arriving at the queue.

5.1.1 *Effect of* ϕ *on the good-state stationary measure* This section illustrates²⁹ how the stationary measure at a good server varies with the inverse average service rate, ϕ , defined in (1). We use the shorthand³⁰ $\phi_N^* > 1$ to denote the unique solution to $[\phi/(1+\phi)]^{N+1} = (\phi-1)/(\phi+q)$.

²⁸Where $\Delta(M) := \{x \in \mathbb{R}^{M}_{+} : \sum_{i} x_{i} = 1\}.$

²⁹All numerical illustrations of the stationary measure in this section are for the value q = 1/2. The values of *N* and *M* are chosen for clarity of illustration and are not necessarily equilibrium values.

³⁰Omitting the dependence on q.

Theoretical Economics 14 (2019)



FIGURE 2. The stationary measure of the queue length conditional on the server being in the good state with $\phi < 1$ under perfect revelation (left panel) and imperfect revelation (right panel).



FIGURE 3. The stationary measure of the queue length conditional on the server being in the good state with $\phi > \phi_N^*$ under perfect revelation (left panel) and imperfect revelation (right panel).

Decreasing when $\phi < 1$ In this case, service is faster than arrivals, so shorter queues are more likely than longer ones, as is apparent in Figure 2. The effect of fast service is further exacerbated by the "renewal" effect of the uninformed first in line reneging after *N* unsuccessful service events, causing the entire queue to clear. The stationary distribution therefore exhibits exponential decline under imperfect revelation and for the values n = 1, ..., N under perfect revelation. The jump down between n = N and n = N + 1 occurs because such a transition is possible only if the first in line knows that the server is good, and similarly for the jump between N + 1 and N + 2 when q < 1. For n = N + 2, ..., M, the distribution declines exponentially.

Increasing when $\phi > \phi_N^*$ In this case service is so slow as to outweigh the renewal effect. The stationary measure is therefore increasing over its entire support, as illustrated in Figure 3. Notice that as *M* increases without bounds, y_1 tends to zero. (See Lemma A.1.) Consequently, if long queues are most likely at a good server, then arriving at the first position in line makes an agent almost certain that the server is bad.

On the interval $[1, \phi_N^*]$ When $\phi = 1$, service is exactly as fast as arrivals. If agents never reneged, but waited in line until served, then every queue length up to M would be

equally likely. For N < M, this is in fact the case when $n \ge N + 2$, and the stationary measure is uniform over these values. When $1 \le n \le N$ or when $M \le N$, the renewal effect is a force for emptying the queue, and the stationary measure is linearly decreasing over these values. For $\phi = \phi_N^*$, service is exactly slow enough to just offset the renewal effect. The stationary measures in these knife-edge cases, as well as on $[1, \phi_N^*]$, are discussed in the Supplemental Material. Observe that the interval $[1, \phi_N^*]$ vanishes $(\phi_N^* \to 1)$ as $N \to \infty$.

5.2 Equilibrium posteriors and inference on queue lengths

In this section, we describe the agents' posterior beliefs and inference in equilibrium. We then state a result that describes the relationship between private and social learning: Lemma 3 shows that, at any point in time, an agent later in the queue is more optimistic than those ahead of her, conditional on no agent having yet observed service.

Let $\bar{\mu}_n^t$ denote the belief about the server state formed, in accordance with E.1., by an agent having arrived at the queue at the *n*th position and having observed *t* periods without service. The upper-bar notation emphasizes the dependence on the strategy $\sigma^*(q, N, M)$.³¹

Consider $\bar{\mu}_n^0$, the belief formed by an agent upon arriving at the *n*th position. In equilibrium, the agent believes that the system is in the steady state induced by $\sigma^*(q, N, M)$. Therefore, $\bar{\mu}_n^0$ is based on the state-dependent stationary measures of queue lengths, $(y_n)_{n=1}^{M+1}$ and $(x_n)_{n=1}^{M+1}$, described in Section 5.1. For every $n \le M + 1$, Bayes' rule gives

$$\bar{\mu}_n^0 := \frac{\mu y_n}{\mu y_n + (1 - \mu) x_n}.$$
(13)

For $N < n \le M + 1$, we have $x_n = 0$ so that $\bar{\mu}_n^0 = 1$. For $n \le \min\{N, M\}$, it was argued above that x_n is independent of n. Therefore, $\bar{\mu}_n^0$ depends on n only through y_n . As $\bar{\mu}_n^0$ is increasing in y_n , the results on the form of the stationary measure in the previous section imply that $\bar{\mu}_n^0$ is decreasing, constant, and increasing in n for $\phi < \phi_N^*$, $\phi = \phi_N^*$, and $\phi > \phi_N^*$, respectively.

Now consider $\bar{\mu}_n^t$, the posterior of the agent at the *n*th position, $n \le N$, who has observed *t* unsuccessful service events. It is derived from $\bar{\mu}_n^0$ according to (5).

Recall that we defined $\bar{\eta}_n$ to be the probability that an agent arriving at the queue at the *n*th position, for $n \leq M$, attaches to the first in line being uninformed, conditional on the server being good. In equilibrium, we have

$$\bar{\eta}_n = \frac{(1-\alpha)^{n-1} y_1}{y_n}.$$
(14)

The numerator gives the stationary probability of arriving at the *n*th position behind a first in line who has never observed service, conditional on the server being in the good state. This is the event that an agent joined the queue at the first position in line and subsequently observed n - 1 periods without service.

³¹In this section, to lighten notation, we omit the dependence of the beliefs ($\bar{\mu}_n^t$ and $\bar{\eta}_n$) and of the stationary measures (y_n and x_n) on (q, N, M).

The next lemma compares agents' posterior beliefs along any given queue under the strategy $\sigma^*(q, N, M)$. This comparison is not trivial. Those ahead of the *n*th in line may (for $n \le \min\{N, M\}$ and $\phi < \phi_N^*$) have been more optimistic than her when they joined the system because they arrived at a shorter queue. However, they have been waiting in the queue for longer and unless they have observed service, waiting will have depressed their belief about the server state. We show that, regardless of how much time she has spent in the queue, an agent is always more optimistic than those ahead of her, conditional on no agent having yet observed service.

The intuition for this result follows from the nesting of agents' information partitions. The n + 1th in line has observed strictly less than the *n*th in line, so her beliefs about the server state are an expectation of the *n*th in line's beliefs. This expectation places positive weight on the *n*th in line knowing that the server is good; that is, $\bar{\mu}_{n+1}^{t-1}$ is an average of 1 and $\bar{\mu}_n^t < 1$. Such an average must be above $\bar{\mu}_n^t$. In fact, if one took a snapshot of the posteriors held by the agents in a queue at any calendar date τ , the sequence of posteriors would be the realization of a martingale. The lemma is proved in Appendix A.3.

LEMMA 3. If $\bar{\mu}_n^t < 1$, then $\bar{\mu}_{n+1}^{t-1} > \bar{\mu}_n^t$.

6. Equilibrium existence

6.1 Equilibria with perfect revelation

In this section, we provide a lower bound on δ above which the strategy $\sigma^*(q^*, N^*, M^*)$ satisfies the equilibrium conditions given in Lemma A and, thus, constitutes an equilibrium with perfect revelation; that is, we provide a sequence of lemmas to prove Proposition 1(i).

The next lemma, proved in Appendix A.4, establishes the existence, for every $M \ge 1$, of a number N^* of periods for which the first in line experiments that satisfies the equilibrium condition (A.iii). Observe that this solution may have $N^* < M$ or $N^* \ge M$. Lemma 5 gives sufficient conditions under which $N^* < M$.

LEMMA 4. For every $M \ge 1$ and $(\alpha, \delta, \mu) \in (0, 1)^3$, there exists a pair (q^*, N^*) with $q^* \in (0, 1]$ and $N^* \ge 0$ such that $N^* = \mathcal{N}(1, \bar{\mu}_1^0(q^*, N^*, M))$.

This result hinges on the following observation, formalized in Lemma A.1. Fix M and consider $\bar{\mu}_1^N(1, N, M)$, which is a first in line's belief after N service failures, derived according to Bayes' rule from the stationary measures of queue lengths induced by the strategy $\sigma^*(1, N, M)$. Increasing the duration N of experimentation by other first in lines (thereby changing the state-dependent stationary measures) ultimately depresses $\bar{\mu}_1^N(1, N, M)$. This is the case even though the effect of a higher N on $\bar{\mu}_1^0(1, N, M)$ —a first in line's belief upon joining the queue—might be positive. Thus, enough bad news from private learning ultimately dominates any good news from social learning. We show that, as a result, there exists a pair (q^*, N^*) such that $\bar{\mu}_1^{N^*}(q^*, N^*, M)$ just hits the first in line's threshold μ_1 defined in (10).

For any given (α, δ, μ) and M, the pair (q^*, N^*) is not necessarily unique.³² The possibility of multiple equilibria arises because longer experimentation by other firsts in line can increase $\bar{\mu}_1^0$ —the belief on the good state formed upon arriving at the first position in line—in turn making longer experimentation optimal. Lemma A.1(b) shows that this process cannot continue indefinitely, implying that the set of possible equilibrium values of (q^*, N^*) is bounded. But there is no clear monotonicity that ensures uniqueness.

The next lemma, proved in Appendix A.5, provides a sufficient condition under which the equilibrium conditions (A.i) and (A.iv) are satisfied.

LEMMA 5. Let $M^* = \mathcal{M}$ and $N^* = \mathcal{N}(1, \bar{\mu}_1^0(q^*, N^*, M^*))$. If

$$V_{N^*} > \frac{\alpha^2}{\alpha^2 - \psi(1 - \delta)} > 0,$$
 (15)

then (a) equilibrium condition (A.iv) holds and (b) equilibrium condition (A.i) holds.

The equilibrium condition (A.iv) requires that no agent in the queue wants to autonomously renege before the queue reaches length N^* , the point at which an uninformed first in line reneges, thereby revealing that she is uninformed. Equivalently, the *n*th agent in line must agree to wait for $N^* - n + 1$ periods before she is rewarded with learning the first in line's information. The sufficient condition (15) ensures that the *n*th agent in line does not want to autonomously renege before the queue reaches length N^* , *even if the informational benefit from obtaining the first in line's information*.³³ Equivalently, it says that the *n*th agent in line is willing to experiment for $N^* - n + 1$ periods, based solely on her own observation of the server (in)activity and on her belief $\bar{\mu}_n^0(q^*, N^*, M^*)$ formed upon joining the queue. It follows that (15) is sufficient for the agents at positions 2, ..., N^* to want to herd on the first in line and, therefore, for equilibrium condition (A.iv) to hold. Finally, we show that when (15) holds, we have $N^* < M^*$ so that condition (A.i) is satisfied.

Lemma 6 ensures that for any w > 1, α , and μ , there exists a critical value of δ such that (15) holds for all discount factors above the critical value. This is the case, in particular, when δ approaches unity so that congestion costs vanish, even though a higher discount factor can mean that an uninformed first in line experiments for longer. The set of parameters satisfying Lemma 6 is illustrated in Figure 4. The lemma is proved in Appendix A.6.

LEMMA 6. Let $M^* = \mathcal{M}$ and $N^* = \mathcal{N}(1, \bar{\mu}_1^0(q^*, N^*, M^*))$. For every $(\alpha, \mu) \in (0, 1)^2$, (15) is satisfied when

$$\delta > 1 + \frac{1}{e \ln(1-\alpha)} \frac{1}{(1-\alpha)^2} \frac{\alpha w \mu}{1-\mu},$$
(16)

³²Figure 1 features several examples of multiple equilibria.

³³Focussing on sufficient conditions that ignore the informational benefit derived from observing the first in line allows us to make an argument that does not depend on the calculated values of the stationary distribution in the good state, but instead relies on bounds on that stationary distribution.

Theoretical Economics 14 (2019)

$$\delta > 1 - \alpha^2,\tag{17}$$

$$\delta > \delta_0(\alpha, \mu),\tag{18}$$

where $\delta_0(\alpha, \mu)$ is the unique value of $\delta \in (0, 1)$ satisfying

$$\ln w = (1 - \delta) \left(4\phi + \frac{\ln(1 - \delta) - \ln \frac{w\alpha\mu}{1 - \mu} + \ln \left[-\frac{e}{2}\ln(1 - \alpha) \right]}{\frac{1}{2}\ln(1 - \alpha)} \phi + \frac{1}{\alpha^3(1 - \alpha)} \right) - \ln \delta.$$



FIGURE 4. The shaded region represents the set of parameter values $(\alpha, \delta) \in (0, 1)^2$ that satisfy Lemma 6. The dotted line represents condition (16), the dashed line represents condition (17), and the solid line represents condition (18). We use $\mu = 0.1$, $\mu = 0.5$, and $\mu = 0.9$, from left to right, and w = 4.

In summary, Lemmas 4, 5, and 6 together prove that there exists a critical value of the discount factor δ such that an equilibrium with perfect revelation exists for all discount factors above the critical value. This establishes Proposition 1(i).

6.2 Equilibria with imperfect revelation

In this section, we complete the proof of Proposition 1(ii) and 1(iii). First, in Lemma 7, we establish the existence of an equilibrium with $N^* = M^* = 1$, when $V_1 \ge 1.^{34}$ Then, in Lemma 8, we show that an equilibrium with imperfect revelation cannot exist for high values of δ .

At an equilibrium with $N^* = M^* = 1$ and $q^* = 1$, all agents arrive at the first position in line. This event is, therefore, not informative about the server state, and $\bar{\mu}_1^0(1, 1, 1) = \mu$, so that an agent arriving at the first position decides whether to join the queue based solely on her prior. Given (α, μ) , it is optimal for her to join the queue if and only if she is sufficiently patient. At the same time, she must be sufficiently impatient to find

³⁴Recall that when the value V_1 of being first in line at a server known to be good is less than 1 (the value of the outside option), it is optimal for any agent arriving at the queue to immediately balk and take the outside option, and the unique equilibrium has $M^* = N^* = 0$.

Strategic experimentation in queues 671



FIGURE 5. For each value of α , the light shaded region represents the interval $D_1^*(\alpha, \mu)$, and the dark shaded region the interval $D_2^*(\alpha, \mu)$. We use $\mu = 0.05$, $\mu = 0.5$ and $\mu = 0.95$, from left to right. (Illustrated for w = 4.)

it optimal to renege on the queue after one service failure. Both these conditions hold when $\delta \in D_1^*(\alpha, \mu)$, where

$$D_1^*(\alpha,\mu) := \left(\frac{1}{1+a}, \frac{1}{1+b}\right], \quad \text{with } a = \alpha(w-1)\mu, b = a\frac{1-\alpha}{1-\alpha\mu}.$$
 (19)

At an equilibrium with $N^* = M^* = 1$ and $q^* \in (0, 1)$, the first in line must be indifferent between reneging after $N^* = 1$ service failure and experimenting one more period. This is the case if and only if $\delta = \delta_1(\alpha, \mu, q^*)$, where

$$\delta_1(\alpha,\mu,q) := \frac{1}{1 + a \frac{(1-\alpha)(2-q)}{1-\alpha\mu + (1-\alpha)(1-q)}} \in D_2^*(\alpha,\mu).$$
(20)

For each $\delta \in D_2^*(\alpha, \mu)$, defined in (21) below, there exists a $q^* \in (0, 1)$ that solves this equation and, therefore, supports an equilibrium with $N^* = M^* = 1$. At such an equilibrium, it is possible for an agent to arrive at the second position in line, and the equilibrium requires that this agent balks. Because of the simple form taken by the stationary distributions of queue lengths, the posterior belief $\bar{\mu}_0^2(q^*, 1, 1)$ she forms upon arrival is equal to the posterior belief $\bar{\mu}_1^1(q^*, 1, 1)$ held by the first in line after observing one service failure. The threshold $\underline{\mu}_n$ defined in (10) increases with *n*, so if the first in line is indifferent between reneging after one or two service failures, then balking is strictly optimal for the agent arriving at the second position in line.

Finally, we define

$$D_2^*(\alpha,\mu) := \left(\frac{1}{1+b}, \frac{1}{1+c}\right), \quad \text{with } c = a \frac{1-\alpha}{1-\frac{1}{2}(1+\mu)\alpha}.$$
 (21)

The adjacent intervals $D_1^*(\alpha, \mu)$ and $D_2^*(\alpha, \mu)$ are illustrated for different values of the prior μ in Figure 5. The proof of Lemma 7 is given in Appendix A.8.

LEMMA 7. Fix $(\alpha, \mu) \in (0, 1)^2$. An equilibrium with $M^* = N^* = 1$ and $q^* = 1$ exists if and only if $\delta \in D_1^*(\alpha, \mu)$. An equilibrium with $M^* = N^* = 1$ and $q^* \in (0, 1)$ exists if and only if $\delta \in D_2^*(\alpha, \mu)$; in this case, q^* is the unique solution to $\delta = \delta_1(\alpha, \mu, q^*)$.

Lemma 8 shows that for $\sigma^*(q^*, N^*, M^*)$ to constitute an equilibrium with imperfect revelation, the discount factor δ must be bounded away from 1. This is intuitive: at an equilibrium with imperfect revelation, agents must be impatient enough to refuse to join long queues. The second and third equilibrium conditions in (B.ii) stipulate that the agent arriving at the M^* + 1th position must balk and, in particular, that she does not wish to join the queue and experiment for m = 1 period. However, when the discount factor is close to 1, waiting is virtually costless. A contradiction. Lemma 8 is proved in Appendix A.9.

LEMMA 8. Given any $(\alpha, \mu) \in (0, 1)^2$, there exists $\overline{\delta}(\alpha, \mu) < 1$ such that an equilibrium with imperfect revelation cannot exist for any $\delta \in (\overline{\delta}(\alpha, \mu), 1)$.

7. Other equilibria?

In this section, we briefly discuss whether other symmetric steady-state Bayesian equilibria might exist for this game; we refer the reader to the Supplemental Material for the details. We consider strategies under which the queue length is bounded. Hence, the stationary distributions (with appropriately enlarged state spaces) are well defined and an agent's inference upon arrival at the queue proceeds as in the previous sections. Moreover, since agents' information sets are nested, it remains true that whenever an agent sees her predecessor renege, she also finds it optimal to renege. We therefore consider potential equilibria in the class of "herding" strategies, where later arrivals in the queue herd on the decisions of those ahead of them. These generalize σ^* in a natural way.

Herding strategies focus on particular agents at particular positions in the queue whose actions are informative. We call these agents *herding leaders*. The strategy of a herding leader is to pick a number of periods to experiment, and to renege if no service is observed before that time has elapsed or if someone ahead of her reneges. The strategy of a *herding follower* is to focus on the closest herding leader ahead of her in the queue, and to renege if and only if she does. Once in the queue, a herding leader's decision to renege depends on her private learning and the publicly observable actions of other leaders. A herding follower's decision to renege depends only on the publicly observable herding leaders' actions.

In Appendix B.4 in the Supplemental Material, we show that any given strategy with more than one herding leader cannot sustain an equilibrium when $\alpha \ge 1/2$ and δ is sufficiently large. This result rests on the following observation. If a herding leader is not the first in line, then the agent just ahead of her in the queue must not want to renege when she does. However, there is a chance that the agent just ahead of her has been waiting longer than one period without observing service (because a previous intermediate herding leader just reneged). At these histories, the intermediate herding leader

is much more optimistic about the server state than the agent just ahead of her. When $\alpha \ge 1/2$ and δ is sufficiently large, this effect is so strong that the intermediate herding leader does not wish to renege if the agent just ahead of her prefers to remain in line, a contradiction.

8. The effects of social learning in queues

The interplay of private and social learning depends on whether the good server is fast (high α) or slow. We show that although a larger α encourages individual learning, it is bad for social learning.

As illustrated in Figures 2 and 3, a new arrival is more likely to find an empty queue at a good server when the service rate is high. Thus, there exists a threshold service rate above which arriving at the first position in line is good news $(\bar{\mu}_1^0(q^*, N^*, M^*) > \mu)$ in equilibrium and the first in line experiments for longer than she would as a single decision-maker with prior μ . Below that threshold, the converse is true. Although the value of this threshold varies with N^* and M^* , we construct bounds on the threshold that are independent of these parameters. The proof of Lemma 9 is given in Appendix A.10.

LEMMA 9. If $\alpha > 2/3$, then every equilibrium with $N^* > 1$ has the property that $\mu < \bar{\mu}_1^0(q^*, N^*, M^*)$, so being first in line is good news about the state. If $\alpha < 2/(3 + \sqrt{5})$, then every equilibrium with $M^* > N^* > 1$ has the property that $\mu > \bar{\mu}_1^0(q^*, N^*, M^*)$ and so being first in line is bad news about the state.

When δ approaches 1, the effect of α on the first in line's equilibrium willingness to experiment is more dramatic. There are two effects that interact. The decreasing costs of delay incline the first in line to experiment longer. This in turn affects the stationary distributions in both states and the posterior $\bar{\mu}_1^0(q^*, N^*, M^*)$ of the first in line. This posterior might decrease, curtailing her equilibrium willingness to experiment. For $\alpha > 1/2$, the first effect dominates and, as δ approaches 1, the first in line's willingness to experiment grows without bounds. For $\alpha < 1/2$, the second effect dominates, so that arriving at the first position makes an agent almost certain that the server is bad and we have $N^* = 1.^{35}$ For $\alpha = 1/2$, the effects balance out and N^* tends to a finite constant, $c(\mu)$, that increases with the prior μ . Lemma 10 is proved in Appendix A.11.

LEMMA 10. For every $(\alpha, \mu) \in (0, 1)^2$, as $\delta \to 1$, N^* converges to

- (a) $+\infty$ when $\alpha > 1/2$
- (*b*) 1 *when* $\alpha < 1/2$
- (c) $c(\mu)$, such that $1 < c(\mu) < \infty$, when $\alpha = 1/2$.

It is worth noting that when $\alpha < 1/2$, there *cannot* be a pure strategy equilibrium as $\delta \rightarrow 1$. Because $N^* = 1$, if $q^* = 1$, then all agents arrive at an empty server regardless of the state ($y_1 = x_1 = 1$), and this event is uninformative about the server state so

³⁵In that case, equilibrium requires $q^* < 1$.

that $\bar{\mu}_1^0(q^*, N^*, M^*) = \mu$. Clearly then, we could not have $N^* = 1$ for every $\mu \in (0, 1)$. This contradiction is resolved if $q^* \in (0, 1)$. In this case, at a good server, there is a positive probability that the agent joins the queue at the second position and becomes an informed first in line if her predecessor is served. From there on, the queue can grow infinitely long. Since the server is slow, it then takes a very long time on average for that queue to clear. Hence, arriving at an empty queue is exceedingly unlikely, and in that event an agent infers that the server is almost certainly bad and chooses $N^* = 1$.

In our equilibrium, once one agent has observed service, no agent will renege on the queue until the queue has emptied. This implies that an entire queue being served or reneging en masse is bad for social learning, as subsequent arrivals have to re-learn what past generations already knew. The expected time between two empty queues is called the *busy period*. In our model, it is a measure of how an equilibrium allows information to propagate through time, which we term social memory.

By standard results for positive recurrent Markov processes, the mean return time to the state in which an agent arrives at the queue in first position is given by $1/y_1$ at a good server.³⁶ Thus, when y_1 is small in equilibrium and queues tend to fill up, the busy period can be very long, which is good for preserving what was already learnt. Conversely, when y_1 is large in equilibrium and queues tend to clear, the busy period is short and the social memory is frequently reset. For instance, the equilibria described in Lemma 10(b) have the property that $y_1 \rightarrow 0$ (as $\delta \rightarrow 1$) so that social memory is very good; in contrast the amount of equilibrium experimentation is small and $N^* = 1$.

9. CONCLUSIONS AND FURTHER WORK

Some of our motivating examples suggest an alternative modelling choice for the queue discipline. For instance, the first-come–first-served discipline could be replaced by an egalitarian random order processing discipline where existing service capacity is allocated equiprobably to all agents currently in the queue.³⁷ In such a model, when an agent reneges, the congestion faced by the remaining agents is lessened, offsetting their increased pessimism about the server state. Consequently, if the first agent in line reneges on the queue, then even though all agents behind her adopt her belief, it is no longer the case that all of them wish to renege with her. This significantly complicates the stochastic process governing the state-dependent steady-state distributions of queue lengths. In particular, the system of dynamic equations produced does not admit a closed-form solution.

Nevertheless, issues similar to those in our model arise, and we hope that the results derived here will be useful in analyzing these related problems. One important feature of the queuing structure is that the individuals' information is nested: any individual has collected strictly less information than those ahead of her in the queue. The fundamental insight our analysis offers to the more general question of experimentation with informational and payoff externalities is that strategy profiles in which agents con-

³⁶It is $N^* + 1 - q^* = 1/x_1$ at a bad server. (See, for example, Brémaud (1999, p. 104).)

³⁷That is, if at the beginning of period τ there are n_{τ} agents in line and service capacity $g_{\tau} < n_{\tau}$ is produced, then each agent is served with probability g_{τ}/n_{τ} .

centrate the social learning on certain focal agents might also result in such nesting of information, and these profiles are good candidate equilibria in more general settings.

The main takeaway for the design and management of queuing systems is that when there is uncertainty about the service rate, the interplay of private and social learning means that queue lengths are more variable at observable FCFS queues than if the queue were not observable. The extra variability comes from the mass exits that happen cyclically at a bad server, but can also happen at a good server.

Such dramatic events are rarely observed in real life. Our model ignores certain frictions that might mitigate the frequency and size of these mass exits. First we assume that there is no service at all at a bad server. If instead we had assumed that in the bad state, the server did serve, but at a slower rate than a good server, then observing one agent renege could still trigger a herd, but need not lead to the entire queue clearing. The same would be true if we had assumed that agents have different discount factors or different opportunity costs of queueing. And sometimes we are just loath to trust the judgement of others and will try to operate an ATM or a parking meter even after being told that it is broken.

Appendix A

A.1 Proof of Lemma 2

Taking a difference and substituting for ψ gives

$$U_n(m+1,\mu_n^0) - U_n(m,\mu_n^0) = \frac{\mu_n^0 \alpha}{\psi} \delta^m (1-\alpha)^m \bigg\{ \psi^n \delta w - 1 - \frac{(1-\mu_n^0)\psi(1-\delta)}{\mu_n^0 \alpha (1-\alpha)^m} \bigg\}.$$
 (A.1)

The term in braces is strictly decreasing in *m* and tends to negative infinity as $m \to \infty$. The function $U_n(\cdot, \mu_n^0)$ is, therefore, strictly quasiconcave in *m* and has a maximal value on $m \in \mathbb{N}^0$. Thus, there is a solution to the problem $\max_{m \in \mathbb{N}^0} U_n(m, \mu_n^0)$. The maximizing *m* is described by the smallest *m* for which $U_n(m+1, \mu_n^0) - U_n(m, \mu_n^0)$ is nonpositive. The solution is generically unique by the strict monotonicity of the braces in (A.1). Setting the braces in (A.1) to equal zero allows us to determine m^* in (9).

Using the second expression in (5) and setting the braces in (A.1) equal to zero gives

us the expression in (10) for $\underline{\mu}_n$, the *n*th in line's cutoff posterior. If μ_n^0 is such that $U_n(m^* + 1, \mu_n^0) = U_n(m^*, \mu_n^0)$ or, equivalently, $\mu_n^{m^*} = \underline{\mu}_n$, then it is optimal to experiment for m^* periods or $m^* + 1$, or to randomize between the two.

A.2 Stationary measures of queue lengths

In addition to $(x_n(q, N, M))_{n=1}^{M+1}$ and $(y_n(q, N, M))_{n=1}^{M+1}$, we define $(z_n(q, N, M))_{n=1}^{N+1} \in \Delta(N+1)$ to be the stationary probability under the strategy $\sigma^*(q, N, M)$ of arriving at the *n*th position behind a first in line who has never observed service, conditional on the server being in the good state.³⁸ Proposition A.1 gives exact expressions for y_n , z_n , and x_n . In the statement of the proposition there are two excluded values of α : the first

³⁸The remaining part of the stationary distribution can be found by taking the difference, $y_n - z_n$. Observe that an agent cannot have arrived at the first position and know that the server is good. Therefore the corresponding state has zero measure, and $y_1 = z_1$.

 $\alpha_N^* < 1/2$ is defined as the solution to $\alpha(\phi + q)(1 - \alpha)^{N+1} = 1 - 2\alpha$; the second is $\alpha = \frac{1}{2}$. In both of these cases the stationary measure in the good state exists, but has a different functional form from the one given in Proposition A.1: The stationary measure when $\alpha = \alpha_N^*$, can be obtained by taking the limit as $k_N \to \infty$ of the expressions in Proposition A.1; the stationary measure when $\alpha = \frac{1}{2}$ is described in the Supplemental Material.

PROPOSITION A.1. Assume that $\alpha \notin \{1/2, \alpha_N^*\}$. If all agents follow the strategy $\sigma^*(q, N, M)$, then, conditional on the server being in the good state, the unique stationary measure satisfies $z_n = (1 - \alpha)^{n-1} y_1$ for n = 1, 2, ..., N + 1.

If
$$N < M$$
, then

$$y_{n} = B \cdot \begin{cases} \phi^{n-1} - k_{N}, & n = 1, 2, \dots, N, \\ \phi^{n-1} - k_{N} \frac{q + (1-q)\alpha\phi^{2}}{\alpha\phi^{2} + q(1-\alpha)}, & n = N+1, \\ \phi^{n-1} - k_{N} \frac{1+q\phi}{\phi^{2}(\phi+q)}\phi^{n-N}, & n = N+2, \dots, M+1, \end{cases}$$
(A.2)

where

$$B^{-1} = \frac{1 - \phi^N}{1 - \phi} - Nk_N + \frac{\phi^N - \phi^{M+1}}{1 - \phi} \left(1 - \frac{k_N}{\phi^{N+1}} \frac{1 + q\phi}{\phi + q} \right) + (1 - q)k_N \frac{1 - \phi^2}{\phi(\phi + q)}.$$
 (A.3)

If $M \leq N$, then

$$y_n = B \cdot \begin{cases} \phi^{n-1} - k_N, & n = 1, 2, \dots, M, \\ \phi^M - \frac{k_N}{\phi}, & n = M + 1, \end{cases}$$
(A.4)

where

$$B^{-1} = \frac{1 - \phi^{M+1}}{1 - \phi} - k_N (M + \phi^{-1}).$$
(A.5)

In both cases, $k_N := \alpha(\phi + q)(1 - \alpha)^{N+1} / [\alpha(\phi + q)(1 - \alpha)^{N+1} + 2\alpha - 1].$

Conditional on the server being in the bad state, the unique stationary measure satisfies, for N < M,

$$x_n = \begin{cases} \frac{1}{N+1-q}, & n = 1, 2, \dots, N, \\ \frac{1-q}{N+1-q}, & n = N+1, \\ 0, & n = N+2, \dots, M+1, \end{cases}$$
(A.6)

and for $M \leq N$,

$$x_n = \begin{cases} \frac{1}{N+1-q}, & n = 1, 2, \dots, M, \\ \frac{N-M+1-q}{N+1-q}, & n = M+1. \end{cases}$$
(A.7)

PROOF. Situation I. Good server under perfect revelation. We begin by considering the recursions that the stationary distribution of the queue lengths must satisfy when N < M and $q \in (0, 1]$.

Consider first the state in which the queue length is n = 1. It is possible to enter this state if there were previously r agents in line and a service capacity of at least r was produced (probability α^r). It is also possible to enter state n = 1 if there were N or N + 1agents in line in the previous period and the first in line had never observed service, was not served, and reneged, causing the entire queue to renege. Thus we can write

$$y_1 = z_N(1-\alpha)(1-\alpha(1-q)) + \sum_{r=1}^M \alpha^r y_r + \alpha^M y_{M+1},$$

where z_N is the stationary probability of a queue length N with an uninformed first in line. The last term arises because there are M agents in line both in state y_M and in state y_{M+1} .

For n > 1, $n \neq N + 1$ and n < M the queue can enter state *n* if no service occurred last period (probability $1 - \alpha$) and there were n - 1 agents in the line, or if r - (n - 1) agents are served (probability $(1 - \alpha)\alpha^{r-n+1}$) and the queue was previously in state *r*. Thus

$$y_n = (1 - \alpha) \sum_{r=n-1}^{M} \alpha^{r-n+1} y_r + (1 - \alpha) \alpha^{M-n+1} y_{M+1}.$$

The system transits to the state where the queue length is N + 1 if the queue is length N, there is no service, and either (a) the first in line knows that the server is in the good state or (b) the first in line is uninformed but her randomizing determines that she wait one more period (probability 1 - q). A second route to entering state N + 1 is if the queue was previously in state r > N and exactly r - N agents were served. Hence

$$y_{N+1} = (1-\alpha)(y_N - z_N) + (1-\alpha)(1-q)z_N + (1-\alpha)\sum_{r=N+1}^M \alpha^{r-N}y_r + (1-\alpha)\alpha^{M-N}y_{M+1}.$$

A little rearranging gives

$$y_{N+1} + q(1-\alpha)z_N = (1-\alpha)\sum_{r=N}^M \alpha^{r-N}y_r + (1-\alpha)\alpha^{M-N}y_{M+1}.$$

A similar calculation for queues of length N + 2 gives

$$y_{N+2} = (1-\alpha) \left(y_{N+1} - z_N (1-\alpha) (1-q) \right) + (1-\alpha) \sum_{r=N+2}^{M} \alpha^{r-N-1} y_r + (1-\alpha) \alpha^{M-N+1} y_{M+1}$$

or

$$y_{N+2} + (1-q)(1-\alpha)^2 z_N = (1-\alpha) \sum_{r=N+1}^M \alpha^{r-N-1} y_r + (1-\alpha) \alpha^{M-N+1} y_{M+1}.$$

The probability that the queue is of length M equals $y_M + y_{M+1}$, which is the probability that the latest agent arrives at the Mth position and joins the queue or at the M + 1th position and balks. An agent arrives at the Mth position if the queue was of length M – 1 at the end of the last period and no service occurred or it was of length M and exactly one service event occurred:

$$y_M = (1 - \alpha)y_{M-1} + (1 - \alpha)\alpha[y_M + y_{M+1}].$$

An agent arrives at the M + 1th position if the queue was of length M at the end of the last period and no service occurred:

$$y_{M+1} = (1 - \alpha)[y_M + y_{M+1}].$$

Rearranging this gives $y_{M+1} = y_M(1 - \alpha)/\alpha$ and a substitution gives

$$y_M = (1 - \alpha)y_{M-1} + \alpha y_{M+1}$$

This completes our description of the recursion satisfied by the state probabilities $\{y_n\}_{n=1}^M$. It is summarized as

$$y_{n} = \begin{cases} \sum_{r=1}^{M} \alpha^{r} y_{r} + \alpha^{M} y_{M+1} + z_{N} (1-\alpha) (1-\alpha(1-q)), & n = 1, \\ (1-\alpha) \sum_{r=n-1}^{M} \alpha^{r-n+1} y_{r} + (1-\alpha) \alpha^{M-n+1} y_{M+1}, & 1 < n \le N, \\ (1-\alpha) \sum_{r=N}^{M} \alpha^{r-N} y_{r} + (1-\alpha) \alpha^{M-N} y_{M+1} - q(1-\alpha) z_{N}, & n = N+1, \\ (1-\alpha) \left[\sum_{r=N+1}^{M} \alpha^{r-N-1} y_{r} + \alpha^{M-N-1} y_{M+1} - (1-q)(1-\alpha) z_{N} \right], & n = N+2, \end{cases}$$
(A.8)
$$(1-\alpha) \sum_{r=n-1}^{M} \alpha^{r-n+1} y_{r} + (1-\alpha) \alpha^{M-n+1} y_{M+1}, & N+2 < n < M, \\ (1-\alpha) y_{M-1} + \alpha y_{M+1}, & n = M, \end{cases}$$

$$(1-\alpha)\alpha^{-1}y_M, \qquad n=M+1.$$

Any nonnegative solution to this system satisfying $\sum_{n=1}^{M+1} y_n = 1$ is a stationary distribution.

Before solving this system we determine the value of z_N , which is the stationary probability of a queue of length N with an uninformed first in line. Because at any date τ the arrival stage follows both the service and exit stages, if an agent arrives in the queue at the first position at date τ , it must be the case that the agent is uninformed: she arrives *after* the last service stage, and after the exit stage at which a queue of length N or N + 1 would have reneged. Therefore $y_1 = z_1$. The probability that an agent who arrived at the first position in the queue is still not served after N - 1 further arrivals is $(1 - \alpha)^{N-1}$.

Therefore, the stationary probability of a queue length *N* with an uninformed first agent is $(1 - \alpha)^{N-1}y_1$. Following the same argument for queue lengths $n \le N$, we conclude that

$$z_n = (1 - \alpha)^{n-1} y_1, \quad n = 1, 2, \dots, N.$$
 (A.9)

It is now clear that the system (A.8) is homogenous of degree 1.

Let us use the fact that

$$(1-\alpha)\sum_{r=n-1}^{M} \alpha^{r-n+1} y_r = (1-\alpha)y_{n-1} + \alpha(1-\alpha)\sum_{r=n}^{M} \alpha^{r-n} y_r$$

to simplify (A.8):

$$y_{n} = \begin{cases} \alpha(1-\alpha)^{-1}y_{2} + z_{N}(1-\alpha)(1-\alpha(1-q)), & n = 1, \\ (1-\alpha)y_{n-1} + \alpha y_{n+1}, & 1 < n < N, \\ (1-\alpha)y_{N-1} + \alpha y_{N+1} + \alpha(1-\alpha)qz_{N}, & n = N, \\ (1-\alpha)y_{N} + \alpha y_{N+2} - (1-\alpha)qz_{N} + \alpha(1-\alpha)^{2}(1-q)z_{N}, & n = N + 1, \\ (1-\alpha)y_{N+1} + \alpha y_{N+3} - (1-\alpha)^{2}(1-q)z_{N}, & n = N + 2, \\ (1-\alpha)y_{n-1} + \alpha y_{n+1}, & N + 2 < n < M + 1, \\ (1-\alpha)\alpha^{-1}y_{M}, & n = M + 1. \end{cases}$$
(A.10)

We now solve this difference equation. For n = 1, 2, ..., N, we have a difference equation of the form $0 = (1 - \alpha)y_{n-1} - y_n + \alpha y_{n+1}$, with the initial and terminal conditions given, respectively, by the expressions for y_1 and y_N in (A.10). The characteristic polynomial for this difference equation is $(x - 1)(x - (1 - \alpha)/\alpha)$. For $\alpha \neq 1/2$, it admits two distinct roots and the difference equation admits the general solution

$$y_n = K + H\phi^n, \quad \phi := \frac{1-\alpha}{\alpha},$$

where *K* and *H* are arbitrary constants.

Imposing the initial condition on this equation allows us to solve for K and gives

$$y_n = \frac{(1-\alpha)^2 (1-\alpha+q\alpha) z_N}{1-2\alpha} + H\phi^n, \quad n = 1, 2, ..., N.$$

Substituting this into the equations above for y_N , y_{N+1} , and y_{N+2} then gives

$$y_{N+1} = H\phi^{N+1} + \frac{(1-\alpha)^2}{1-2\alpha} z_N [(1-\alpha)(1-q) + (q/\phi)],$$

$$y_{N+2} = H\phi^{N+2} + \frac{(1-\alpha)^2}{1-2\alpha} z_N [\alpha + q(1-\alpha)],$$

$$y_{N+3} = H\phi^{N+3} + \frac{\phi(1-\alpha)^2}{1-2\alpha} z_N [\alpha + q(1-\alpha)].$$

Now let us turn to states $N + 2 < n \le M + 1$. Taking the terminal condition given by the expression for y_M and y_{M+1} in (A.10) and substituting into the y_{M-1} equation gives $y_{M-1} = (\frac{\alpha}{1-\alpha})^2 y_{M+1}$. Hence, $y_n = (\alpha/(1-\alpha))^{M+1-n} y_{M+1}$. Alternatively,

$$y_n = \phi^{n-N-2} y_{N+2}, \quad n = N+2, \dots, M+1.$$

Combining the two parts of the solution, we get

$$y_n = \begin{cases} \frac{(1-\alpha)^2(1-\alpha+\alpha q)z_N}{1-2\alpha} + H\phi^n, & n = 1, 2, \dots, N, \\ \frac{(1-\alpha)^2 z_N}{1-2\alpha} ((1-\alpha)(1-q) + (q/\phi)) + H\phi^{N+1}, & n = N+1, \\ \frac{(1-\alpha)^2 (\alpha+q(1-\alpha))z_N}{1-2\alpha} \phi^{n-N-2} + H\phi^n, & n = N+2, \dots, M+1 \end{cases}$$

We now substitute the value of z_N into the y_1 equation. A rewriting of (A.9) gives

$$z_N = (1 - \alpha)^{N-1} \left(\frac{(1 - \alpha)^2 z_N}{1 - 2\alpha} (1 - \alpha + \alpha q) + H\phi \right).$$

Hence,

$$\frac{(1-\alpha)^2 z_N}{1-2\alpha} (1-\alpha+\alpha q) = \frac{H(1-\alpha)^{N+1} \phi(1-\alpha+\alpha q)}{\left(1-2\alpha-(1-\alpha+\alpha q)(1-\alpha)^{N+1}\right)} = -H\phi k_N,$$

where $k_N := (1 - \alpha + \alpha q)(1 - \alpha)^{N+1} / [(1 - \alpha + \alpha q)(1 - \alpha)^{N+1} + 2\alpha - 1]$ (k_N is defined by our assumption in the statement of the lemma). Substituting into the above then gives

$$y_{n} = H \begin{cases} \phi^{n} - k_{N}\phi, & n = 1, 2, ..., N, \\ \phi^{n} - k_{N}\left(\phi + \frac{q(1-\phi)}{\alpha(\phi+q)}\right), & n = N+1, \\ \phi^{n} - k_{N}'\phi^{n-N-1}, & n = N+2, ..., M+1, \end{cases}$$
(A.11)

where $k'_N := k_N(1 + q\phi)/(\phi + q)$. This gives the final form of the distribution given in the lemma.

To verify that this is a legitimate stationary measure we must check that there exists a scalar *H* such that the y_n , defined by (A.11), are all nonnegative. The terms y_{N+2}, \ldots, y_{M+1} are all proportionate, so it is sufficient to check that y_1, \ldots, y_{N+2} are nonnegative. To address this question we consider three separate cases.

Let α_N^* satisfy $(1 - \alpha + \alpha q)(1 - \alpha)^{N+1} + 2\alpha - 1 = 0$. (Then $\alpha_N^* < 1/2$ and $\alpha_N^* \to 1/2$ as $N \to \infty$.) Furthermore, $k_N < 0$ if $\alpha < \alpha_N^*$ and $k_N > 0$ if $\alpha > \alpha_N^*$. When $\alpha > \alpha_N^*$, k_N is strictly decreasing, with $k_N = 1$ when $\alpha = 1/2$. Furthermore, noticing that for n > 1, $(1 - \alpha)^n + 2\alpha - 1 - \alpha^n$ has three roots on [0, 1] (they are 0, 1/2, and 1) and is strictly convex on (0, 1/2) and strictly concave on (1/2, 1), we obtain that

 $\phi^{N+1} - k_N$ has the same sign as $1 - k_N$. We therefore distinguish the following three cases.

- **Case A.1.1:** $1/2 < \alpha < 1$. Since $0 < k_N < 1$, to ensure $y_1 \ge 0$, we require $H \ge 0$. When H > 0, the terms y_1, \ldots, y_{N+2} decrease (since $\phi < 1$), so it is sufficient to check that $y_{N+2} \ge 0$. This is the case since $\phi^{N+1} \ge k_N$.
- **Case A.1.2:** $\alpha_N^* < \alpha < 1/2$. Since $k_N > 1$, from $y_1 \ge 0$ we must have $H \le 0$. The terms y_1, \ldots, y_N , therefore, decrease (since $\phi > 1$). It is sufficient to check that y_N, y_{N+1} , $y_{N+2} \ge 0$. The first two variables follow from $k_N > \phi^{N+1}$. To verify that $y_{N+2} \ge 0$, full substitution for k is necessary to get an inequality that is linear in q. The two cases q = 0 and q = 1 follow from the above inequalities.
- **Case A.1.3:** $\alpha < \alpha_N^*$. Since $k_N < 1$, from $y_1 \ge 0$ we must have $H \ge 0$. Since $k_N < 0$, all y_n are then positive.

The constant *H* must be chosen so that the y_n defined in (A.11) sum to unity. Thus we choose

$$H^{-1} = \sum_{n=1}^{M+1} \phi^n - (N+1)\phi k_N - k_N \frac{1+q\phi}{\phi+q} \sum_{n=1}^{M-N} \phi^n - k_N \frac{q(1-\phi)}{\alpha(\phi+q)}$$

or

$$H^{-1} = \frac{\phi(1-\phi^{M+1})}{1-\phi} - k_N \phi N - k_N \frac{1+q\phi}{\phi+q} \frac{1-\phi^{M-N+1}}{1-\phi} + (1-q)k_N \frac{(1-\phi^2)}{\phi+q}.$$

It will be convenient to cancel ϕ when we rewrite the above as (A.3) in the lemma. After some algebra it can be verified that for $\alpha \in (0, 1)$, H has the same sign as $1 - k_N$ and we therefore have a legitimate stationary measure with $y_n > 0$ for all $n = 1, \ldots, M + 1$. The uniqueness of this stationary distribution follows from the fact that the strategy described induces an irreducible Markov process on the states $n = 1, \ldots, M$.

Situation II. Good server under imperfect revelation. Now the queue never grows longer than length M, even if the first in line is still experimenting, because no additional agent is willing to join a queue longer than M. The probability of arriving at the M + 1th position (and then balking) depends on whether the first in line is informed. If the first in line is uninformed and there are M agents in line, then N - M further unsuccessful service events occur before the first in line exits or N - M + 1 if she exits after observing N + 1 unsuccessful service event, which her strategy prescribes with probability (1 - q). If the first in line is informed, there can be infinitely many unsuccessful service events. Therefore,

$$y_{M+1} = z_M \sum_{i=1}^{N-M} (1-\alpha)^i + z_M (1-q)(1-\alpha)^{N-M+1} + (y_M - z_M) \sum_{i=1}^{\infty} (1-\alpha)^i.$$

Simplifying yields

$$y_{M+1} = \frac{1-\alpha}{\alpha} (y_M - z_M (1-\alpha)^{N-M} (1-\alpha + q\alpha)).$$
(A.12)

The probability of arriving at the first position equals the probability of a queue of length $1, \ldots, M$ clearing plus the probability of an uninformed first in line reneging after having observed N or N + 1 unsuccessful service events:

$$y_1 = z_1 = \sum_{r=1}^{M} \alpha^r y_r + \alpha^M y_{M+1} + z_1 (1-\alpha)^N (1-\alpha+q\alpha).$$
 (A.13)

For M < N, the probability of arriving at the *n*th position satisfies the same recursion as for $N \le M$,

$$y_n = (1 - \alpha) \sum_{r=n-1}^{M} \alpha^{r-n+1} y_r + (1 - \alpha) \alpha^{M-n+1} y_{M+1}, \quad n = 2, \dots, M,$$
 (A.14)

and the probability of arriving at the *n*th position and the first in line being uninformed is

$$z_n = (1 - \alpha)^{n-1} z_1, \quad n = 2, \dots, M.$$
 (A.15)

The recursion (A.14) gives the same difference equation as before,

$$y_n = \alpha y_{n+1} + (1 - \alpha) y_{n-1}, \quad n = 2, \dots, M,$$

which, for $\alpha \neq 1/2$, admits the same general solution $y_n = K + H\phi^n$ as previously. Rewriting the initial condition (A.13) by substituting (A.14) for y_2 , we obtain

$$y_1 = \frac{\alpha}{1-\alpha} y_2 + z_1 (1-\alpha)^N (1-\alpha+q\alpha).$$

Imposing this on $y_n = K + H\phi^n$, we obtain

$$y_n = \frac{(1-\alpha)^{N+1}(1-\alpha+q\alpha)}{1-2\alpha}z_1 + H\phi^n, \quad n = 1, ..., M.$$

We use $z_1 = y_1$ to solve for *H* in the expression above to obtain

$$y_n = z_1 \frac{\phi^{n-1} - k_N}{1 - k_N}, \quad n = 1, \dots, M,$$

where k_N is as defined previously.

Using this expression for n = M together with the terminal condition (A.12) (where we use (A.15) to simplify z_M) then gives

$$y_{M+1} = z_1 \frac{\phi^M - k_N \phi^{-1}}{1 - k_N}$$

and

$$y_M + y_{M+1} = \frac{z_1}{1-\alpha} \frac{\phi^M - k_N}{1-k_N}.$$

Finally, imposing the condition that $\sum_{n=1}^{M+1} y_n = 1$, we get

$$1 = \frac{z_1}{1 - k_N} \left(\sum_{n=1}^{M-1} (\phi^{n-1} - k_N) + \frac{\phi^M - k_N}{1 - \alpha} \right), \tag{A.16}$$

which simplifies to

$$1 = \frac{z_1}{1 - k_N} \left(\frac{1 - \phi^{M+1}}{1 - \phi} - k_N (M + \phi^{-1}) \right).$$
(A.17)

This determines the last part of the solution.

We now verify that all the y_n are nonnegative. For $2 \le n \le N + 1$, we have that $1 < \phi^{n-1} < \phi^{N+1}$ when $\alpha < 1/2$ and $1 > \phi^{n-1} > \phi^{N+1}$ when $\alpha > 1/2$. So for all admissible values of $\alpha \in (0, 1)$, $\phi^{n-1} - k_N$ lies between $1 - k_N$ and $\phi^{N+1} - k_N$. We have seen in the treatment of $M \le N$ that these two expressions have the same sign for all admissible values of α . It follows that $(\phi^{n-1} - k_N)/(1 - k_N)$ is positive for all admissible $\alpha \in (0, 1)$, so it is sufficient to verify that $z_1 \ge 0$.

From (A.17) we have that $z_1 > 0$ for $\alpha < \alpha_N^*$, because $k_N < 0$. In (A.16) the term in brackets is a sum of positive terms for $\alpha > 1/2$ and a sum of negative terms for $\alpha_N^* < \alpha < 1/2$. It therefore has the same sign as $1 - k_N$ when $\alpha > \alpha_N^*$ and so $z_1 \ge 0$ also for $\alpha > \alpha_N^*$. Hence we have derived a legitimate stationary measure when $M \le N$.

Scenario III: Bad server. We conclude by deriving the stationary distribution conditional on the server being in the bad state. For $M \le N$, the transition equations are $x_1 = \cdots = x_N$ and $x_{N+1} = (1 - q)x_N$. For N < M, they are $x_1 = \cdots = x_M$ and $x_{M+1} = (N - M + 1 - q)x_M$. In each case the result follows from the requirement that the probabilities sum to 1.

A.3 Proof of Lemma 3

If $\bar{\mu}_n^t < 1$ at some time τ , then from (13) and (5) we have that

$$\bar{\mu}_n^t = \frac{\mu y_n (1-\alpha)^t}{\mu y_n (1-\alpha)^t + (1-\mu)x_n}, \qquad \bar{\mu}_{n+1}^{t-1} = \frac{\mu y_{n+1} (1-\alpha)^{t-1}}{\mu y_{n+1} (1-\alpha)^{t-1} + (1-\mu)x_{n+1}}$$

From this it follows that $\bar{\mu}_n^t < \bar{\mu}_{n+1}^{t-1}$ if and only if $(1 - \alpha)y_n/x_n < y_{n+1}/x_{n+1}$. Observe that $x_n \ge x_{n+1}$ by (A.6) and (A.7). Therefore, a sufficient condition for $\bar{\mu}_n^t < \bar{\mu}_{n+1}^{t-1}$ is $(1 - \alpha)y_n < y_{n+1}$.

Among the many ways to arrive at the n + 1th position in line in the good state, one is that the previous agent joined the queue at the *n*th position and there was no service; another is that the previous agent joined the queue at the n + 1th position and exactly one agent was served. We therefore have the bound

$$y_{n+1} \ge (1-\alpha)y_n + \alpha(1-\alpha)y_{n+1}.$$

Rearranging gives

$$y_{n+1} \ge y_n \frac{1-\alpha}{1-\alpha(1-\alpha)},$$
 (A.18)

implying that $y_{n+1} > (1 - \alpha)y_n$ (since $\alpha \in (0, 1)$ and $y_n > 0$). Thus, the sufficient condition is satisfied.

A.4 Proof of Lemma 4

A.4.1 *Intermediate results* We begin the proof of Lemma 4 with two intermediate results (when q = 1). Fix M. Lemma A.1(a) shows that as $\sigma^*(q, N, M)$ prescribes that the first in line experiments for longer (as N increases), the probability of arriving at the queue at the first position when the server is good declines. This is because when N increases there is total probability being spread over more queue lengths so that the probability of any one queue length falls.

Lemma A.1(b) shows that a higher N eventually results in a reduction in the first in line's posterior after N unsuccessful service opportunities, $\bar{\mu}_1^N(1, N, M)$. This is the case even though the effect of a higher N on $\bar{\mu}_1^0(1, N, M)$ might be positive. As N increases, there are many things to take into account: the probability of being first in line at a bad server shrinks to zero, but the probability of being first in line does not necessarily vanish if the server is in the good state. Thus as N increases, arriving first in line may become very good news indeed. After waiting N periods without success, however, the posterior of the first in line is revised so far down that her initial optimism is entirely depleted. The effect of private learning eventually dominates the effect of social learning.

LEMMA A.1. Fix $(\alpha, \delta, \mu) \in (0, 1)^3$ and $M \ge 1$.

- (a) The probability $y_1(1, N, M)$ decreases as N increases for all $\alpha \in (0, 1)$.
- (b) The posterior belief $\bar{\mu}_1^N(1, N, M)$ decreases in N for all $N > 1/\alpha$ and tends to zero as N tends to infinity.

PROOF. (a) Fix $M \ge 1$ and q = 1, and suppose first that $M > N \ge 1$. Making the dependence of y_1 on (q, N, M) explicit, a substitution from (A.2) and (A.3) gives

$$\frac{1}{y_{1}(1, N, M)} = N + \frac{\frac{1-\phi^{N}}{1-\phi} - N + \left(1 - \frac{k_{N}}{\phi^{N+1}}\right)\frac{\phi^{N} - \phi^{M+1}}{1-\phi}}{1-k_{N}}$$
$$= N + \frac{\sum_{i=0}^{N-1}(\phi^{i} - 1) + \left(1 - \frac{k_{N}}{\phi^{N+1}}\right)\frac{\phi^{N} - \phi^{M+1}}{1-\phi}}{1-k_{N}}$$
$$= N + \frac{\phi^{N} - \phi^{M+1}}{1-\phi}\left(1 - \frac{1-\phi^{N+1}}{(1-\phi)(1+\phi)^{N}}\right)$$
$$- \sum_{i=0}^{N-1}\frac{1-\phi^{i}}{1-\phi}\left(1 - \phi + \phi\left(\frac{\phi}{1+\phi}\right)^{N}\right).$$
(A.19)

(To get the final line we substitute $k_N|_{q=1} = \phi^{N+1}/(\phi^{N+1} + (1-\phi)(1+\phi)^N)$.) Simplifying further, we get

$$\frac{1}{y_{1}(1,N,M)} = \sum_{i=0}^{N-1} \phi^{i} + \frac{\phi^{N} - \phi^{M+1}}{1 - \phi} \left(1 - \frac{1 - \phi^{N+1}}{(1 - \phi)(1 + \phi)^{N}} \right) - \phi \left(\frac{\phi}{1 + \phi} \right)^{N} \sum_{i=0}^{N-1} \frac{1 - \phi^{i}}{1 - \phi}$$
$$= \sum_{i=0}^{M} \phi^{i} - \frac{\phi^{N} - \phi^{M+1}}{1 - \phi} \frac{1 - \phi^{N+1}}{(1 - \phi)(1 + \phi)^{N}} - \phi \left(\frac{\phi}{1 + \phi} \right)^{N} \sum_{i=0}^{N-1} \frac{1 - \phi^{i}}{1 - \phi}$$
$$= \frac{1 - \phi^{M+1}}{1 - \phi} - \left(\frac{\phi}{1 + \phi} \right)^{N} \left\{ \phi \sum_{i=0}^{N-1} \frac{1 - \phi^{i}}{1 - \phi} + \frac{1 - \phi^{M-N+1}}{1 - \phi} \frac{1 - \phi^{N+1}}{1 - \phi} \right\}. \quad (A.20)$$

We now focus on the term in braces. This equals

$$\begin{split} \frac{1}{1-\phi} \Bigg[\phi N - \sum_{i=1}^{N} \phi^{i} + \left(1 - \phi^{M-N+1}\right) \sum_{i=0}^{N} \phi^{i} \Bigg] &= 1 + \frac{1}{1-\phi} \Bigg[(N+1)\phi - \phi^{M-N+1} \sum_{i=0}^{N} \phi^{i} \Bigg] \\ &= 1 + \phi \sum_{i=0}^{N} \frac{1 - \phi^{M-N+i}}{1-\phi}. \end{split}$$

Hence,

$$\frac{1}{y_1(1,N,M)} = \sum_{i=0}^{M} \phi^i - \left(\frac{\phi}{1+\phi}\right)^N \left\{ 1 + \phi \sum_{i=0}^{N} \frac{1-\phi^{M-N+i}}{1-\phi} \right\}.$$
 (A.21)

We now study how this changes as N increases, so let us write

$$\frac{1}{y_1(1,N,M)} = K_M - \left(\frac{\phi}{1+\phi}\right)^N H_N.$$

Then, for M > N > 1,

$$\frac{1}{y_1(1,N,M)} - \frac{1}{y_1(1,N-1,M)} = \left(\frac{\phi}{1+\phi}\right)^N \left(H_{N-1} - H_N + \frac{1}{\phi}H_{N-1}\right)$$
$$= \left(\frac{\phi}{1+\phi}\right)^N \left(-\phi\frac{1-\phi^{M-N}}{1-\phi} + \frac{1}{\phi}H_{N-1}\right)$$
$$= \left(\frac{\phi}{1+\phi}\right)^N \left(\frac{1+\phi}{\phi} + \sum_{i=1}^{N-1}\frac{1-\phi^{M-(N-1)+i}}{1-\phi}\right) > 0.$$

This establishes our claim for $M > N \ge 1$.

Now fix $M \ge 1$ and q = 1, and suppose that $M \le N$. We have from (A.4) and (A.5) that

$$\frac{1}{y_1(1,N,M)} = \frac{\frac{1-\phi^{M+1}}{1-\phi} - k_N(M+\phi^{-1})}{1-k_N}$$

Theoretical Economics 14 (2019)

$$= M + 1 + \frac{\sum_{i=0}^{M} (\phi^{i} - 1) + k_{N} (1 - \phi^{-1})}{1 - k_{N}}$$

$$= M + 1 + \sum_{i=0}^{M} (\phi^{i} - 1) \frac{\phi^{N+1} + (1 - \phi)(1 + \phi)^{N}}{(1 - \phi)(1 + \phi)^{N}} - \left(\frac{\phi}{1 + \phi}\right)^{N}$$

$$= M + 1 + \sum_{i=0}^{M} (\phi^{i} - 1) - \left(\frac{\phi}{1 + \phi}\right)^{N} \left[1 + \phi \sum_{i=0}^{M} \frac{\phi^{i} - 1}{\phi - 1}\right]$$

$$= \sum_{i=0}^{M} \phi^{i} - \left(\frac{\phi}{1 + \phi}\right)^{N} \left[1 + \phi \sum_{i=0}^{M} \frac{1 - \phi^{i}}{1 - \phi}\right]$$
(A.22)

(to get the final line, we substitute $k_N|_{q=1} = \phi^{N+1}/(\phi^{N+1} + (1 - \phi)(1 + \phi)^N))$ and it is immediate that $y_1(1, N, M)$ decreases as *N* increases.

(b) Fix $M \ge 1$ and q = 1. From (5) and (13),

. .

$$\frac{\bar{\mu}_1^N(1, N, M)}{1 - \bar{\mu}_1^N(1, N, M)} = \frac{\mu}{1 - \mu} y_1(1, N, M) N(1 - \alpha)^N.$$

To show that $\bar{\mu}_1^N(1, N, M)$ decreases in N it is sufficient to show that $N(1-\alpha)^N$ decreases in N, since $y_1(1, N, M)$ decreases in N from part (a). But $N(1-\alpha)^N$ is quasiconcave in N, and decreases in N for all $N > 1/\alpha$. Finally, $N(1-\alpha)^N y_1(1, N, M)$ converges to zero as N increases because $y_1(1, N, M) \le 1$.

A.4.2 *Proof of Lemma* 4 Fix $(\alpha, \delta, \mu) \in (0, 1)^3$ and $M \ge 1$. We show that there exists $N^* \ge 0$ and $q^* \in (0, 1]$ such that for an agent arriving at the first position in line, $\sigma^*(q^*, N^*, M)$ is a best response to all other agents adhering to $\sigma^*(q^*, N^*, M)$.

Suppose that all other agents use the strategy $\sigma^*(q, N, M)$, where M and N are two nonnegative integers and $q \in (0, 1]$. This determines the state-dependent stationary measures of queue lengths and, in particular, $y_1(q, N, M)$. These are used by every agent to update her prior upon arriving at the queue. Consider the belief $\bar{\mu}_1^N(q, N, M)$ of an agent who joined the queue at the first position and observed N failures.

To accommodate mixed strategies, we let $\tilde{N} := N + 1 - q$ (this can be thought of as the expected duration of a first in line's experimentation) and $\tilde{q}(\tilde{N}) := \lfloor \tilde{N} \rfloor + 1 - \tilde{N}$. We wish to study the properties of $\bar{\mu}_1^N(q, N, M)$ as a function of \tilde{N} . We therefore define³⁹

$$\mathcal{G}(\tilde{N}) := \bar{\mu}_1^{\lfloor \tilde{N} \rfloor} \big(\tilde{q}(\tilde{N}), \lfloor \tilde{N} \rfloor, M \big).$$
(A.23)

By (5) and (13), it satisfies

$$\frac{\mathcal{G}(N)}{1-\mathcal{G}(\tilde{N})} = \frac{\mu}{1-\mu} y_1 \big(\tilde{q}(\tilde{N}), \lfloor \tilde{N} \rfloor, M \big) \tilde{N} (1-\alpha)^{\lfloor \tilde{N} \rfloor}.$$
(A.24)

³⁹Since *M* is fixed, we omit the dependence on *M* so as to lighten notation.

Consider first the term $y_1(\tilde{q}(\tilde{N}), \lfloor \tilde{N} \rfloor, M)$. When $q \in (0, 1]$ and $1 \le N < M$, the analog of (A.20) is

$$\frac{1}{y_1(q,N,M)} = \sum_{i=0}^M \phi^i - \left(\frac{\phi}{1+\phi}\right)^N \left\{ q - 1 + \frac{q+\phi}{1+\phi} \phi \sum_{i=0}^{N-1} \frac{1-\phi^i}{1-\phi} + \frac{1-\phi^{M-N+1}}{1-\phi} \frac{1-\phi^{N+2}+q\phi(1-\phi^N)}{1-\phi^2} \right\},$$
(A.25)

and when $N \ge M$, the analog of (A.22) is

$$\frac{1}{y_1(q,N,M)} = \sum_{i=0}^M \phi^i - \left(\frac{\phi}{1+\phi}\right)^N \left\{ \frac{q+\phi}{1+\phi} \left[1+\phi \sum_{i=0}^M \frac{1-\phi^i}{1-\phi} \right] \right\}.$$
 (A.26)

In both (A.25) and (A.26), the term in braces is strictly increasing in q. Consequently, $y_1(q, N, M)$ is strictly increasing in q for any given N and M. For $1 \le N < M$, taking the limit of the term in braces in (A.25) as $q \to 0$ gives

$$\begin{split} \frac{\phi}{1+\phi} & \left[\phi \sum_{i=0}^{N-1} \frac{1-\phi^i}{1-\phi} + \frac{1-\phi^{M-N+1}}{1-\phi} \frac{1-\phi^{N+2}}{1-\phi} \frac{1}{\phi} - \frac{1+\phi}{\phi} \right] \\ & = \frac{\phi}{1+\phi} \left[\phi \sum_{i=0}^{N-1} \frac{1-\phi^i}{1-\phi} + \frac{1}{1-\phi} \frac{1}{\phi} - \frac{\phi^{N+1}}{1-\phi} + \frac{1-\phi^{M-N}}{1-\phi} \frac{1-\phi^{N+2}}{1-\phi} - \frac{1+\phi}{\phi} \right] \\ & = \frac{\phi}{1+\phi} \left[\phi \sum_{i=0}^{N-1} \frac{1-\phi^i}{1-\phi} + \phi \frac{1-\phi^N}{1-\phi} + \frac{1-\phi^{M-N}}{1-\phi} \frac{1-\phi^{N+2}}{1-\phi} \right], \end{split}$$

so that we have

$$\lim_{q \to 0} \frac{1}{y_1(q, N, M)} = \frac{1 - \phi^{M+1}}{1 - \phi} - \left(\frac{\phi}{1 + \phi}\right)^{N+1} \left\{ \phi \sum_{i=0}^N \frac{1 - \phi^i}{1 - \phi} + \frac{1 - \phi^{M-N}}{1 - \phi} \frac{1 - \phi^{N+2}}{1 - \phi} \right\}$$
$$= \frac{1}{y_1(1, N+1, M)},$$
(A.27)

where the last equality follows from (A.20). For $N \ge M$, taking the limit of the brackets in (A.26) as $q \rightarrow 0$, it is immediate that

$$\lim_{q \to 0} \frac{1}{y_1(q, N, M)} = \sum_{i=0}^M \phi^i - \left(\frac{\phi}{1+\phi}\right)^{N+1} \left[1 + \phi \sum_{i=0}^M \frac{1-\phi^i}{1-\phi}\right]$$
$$= \frac{1}{y_1(1, N+1, M)},$$
(A.28)

where the last equality follows from (A.22). Finally, evaluating (A.21) and (A.22) at N = M, we find that they are equal. Hence, for any given M, $y_1(\tilde{q}(\tilde{N}), \lfloor \tilde{N} \rfloor, M)$ is continuous and decreasing in \tilde{N} for every $\tilde{N} \ge 1$.

In contrast,

$$\begin{split} \lim_{q \to 0} \tilde{N}(1-\alpha)^{\lfloor \tilde{N} \rfloor} &= (\tilde{N}+1)(1-\alpha)^{\lfloor \tilde{N} \rfloor} \\ &> (\tilde{N}+1)(1-\alpha)^{\lfloor \tilde{N}+1 \rfloor} \end{split}$$

Thus, the term $\tilde{N}(1-\alpha)^{\lfloor \tilde{N} \rfloor}$ in (A.24) is increasing at non-integer values of \tilde{N} and has downward jumps at integer values of \tilde{N} .

From (A.24) we therefore have that $\mathcal{G}(\tilde{N})$ is continuous in \tilde{N} at non-integer values of \tilde{N} for every $\tilde{N} \ge 1$ and has downward discontinuities at every integer value of \tilde{N} for $\tilde{N} > 1$. By Lemma A.1(b), $\mathcal{G}(\lfloor \tilde{N} \rfloor) > \mathcal{G}(\lfloor \tilde{N} + 1 \rfloor)$ for every $\tilde{N} > 1/\alpha$, and $\mathcal{G}(\lfloor \tilde{N} \rfloor)$ tends to zero as $\tilde{N} \to \infty$. Moreover, for $\tilde{N} = 0$ and $\tilde{N} = 1$, all agents arrive at the first position in line so that this event is entirely uninformative about the server state and in both cases we have that $\mathcal{G}(\tilde{N}) = \mu$, where μ is the agent's prior about the server state.

Consequently, for every $x \in (0, 1)$, there exists a value $\tilde{N}^*(x) \ge 0$ defined to be the smallest value $\tilde{N} \ge 0$ such that $\mathcal{G}(\tilde{N}) \le x$.

We now consider $\tilde{N}^*(\underline{\mu}_1)$, where $\underline{\mu}_1$ is the first in line's threshold belief defined in (10). For $\mu > \underline{\mu}_1$, $\tilde{N}^*(\underline{\mu}_1)$ may or may not be an integer. We consider each case separately. (If $\mu \le \mu_1$, immediately balking is a best response.)

Case 1. If $\tilde{N}^*(\underline{\mu}_1)$ is a non-integer value, then $\mathcal{G}(\tilde{N}^*(\underline{\mu}_1)) = \underline{\mu}_1$, and setting $N^* = \lfloor \tilde{N}^*(\mu_1) \rfloor$ and $q^* = \tilde{q}(\tilde{N}^*(\mu_1))$ in (A.23) gives

$$\bar{\mu}_1^{N^*}(q^*, N^*, M) = \underline{\mu}_1.$$

Moreover, by (5) and (13), it follows that for every $t = 0, ..., N^* - 1$,

$$\bar{\mu}_1^t(q^*, N^*, M) > \underline{\mu}_1.$$

Consequently, by Lemma 2, when the other agents adhere to $\sigma^*(q^*, N^*, M)$, an agent arriving at the first position finds it optimal to join the queue, to continue at the first $N^* - 1$ failures, and to be indifferent between reneging after N^* or $N^* + 1$ failures. Hence, $\sigma^*(q^*, N^*, M)$ is a best response for the first in line.

Case 2. If $\tilde{N}^*(\mu_1)$ is an integer value, then

$$\mathcal{G}\big(\tilde{N}^*(\underline{\mu}_1)\big) \leq \underline{\mu}_1 < \lim_{\tilde{N} \uparrow \tilde{N}^*(\underline{\mu}_1)} \mathcal{G}(\tilde{N}).$$

Setting $N^* = \lfloor \tilde{N}^*(\mu_1) \rfloor$ and $q^* = 1$ in (A.23), the first inequality above gives

$$\bar{\mu}_1^{N^*}(1, N^*, M) \le \mu_1,$$
 (A.29)

while the second inequality implies that

$$\bar{\mu}_1^t(1, N^*, M) > \mu_1, \quad t = 0, \dots, N^* - 1.$$
 (A.30)

Indeed, $\lim_{\tilde{N}\uparrow N^*} \mathcal{G}(\tilde{N}) \leq \bar{\mu}_1^t(1, N^*, M)$ for every $t = 0, \dots, N^* - 1$, since

$$\lim_{\tilde{N}\uparrow N^*} \frac{\mathcal{G}(N)}{1 - \mathcal{G}(\tilde{N})} = \frac{\mu}{1 - \mu} \lim_{q \to 0} y_1(q, N^* - 1, M) N^* (1 - \alpha)^{N^* - 1}$$
$$= \frac{\mu}{1 - \mu} y_1(1, N^*, M) N^* (1 - \alpha)^{N^* - 1},$$

where the second equality follows from (A.27) and (A.28), while by (5) and (13), for every $t = 0, ..., N^* - 1$,

$$\frac{\bar{\mu}_1^t(1, N^*, M)}{1 - \bar{\mu}_1^t(1, N^*, M)} = \frac{\mu}{1 - \mu} y_1(1, N^*, M) N^* (1 - \alpha)^t.$$

By Lemma 2, it follows from (A.29) and (A.30) that when the other agents adhere to $\sigma^*(q^*, N^*, M)$, an agent arriving at the first position finds it optimal to join the queue, to continue at the first $N^* - 1$ failures, and to renege after the *N*th failure. Hence, $\sigma^*(q^*, N^*, M)$ is a best response for the first in line.

This establishes the existence, for every $M \ge 1$ and (α, δ, μ) of $q^* \in (0, 1]$ and $N^* \ge 0$ satisfying (A.iii). Observe that this solution may have $N^* < M$ or $N^* \ge M$. Lemma 5 gives sufficient conditions under which $N^* < M^*$.

A.5 Proof of Lemma 5

(a) Let $M^* = \mathcal{M}$ (condition (A.ii)) and $N^* = \mathcal{N}(1, \bar{\mu}_1^0(q^*, N^*, M^*))$ (condition (A.iii)). Assume that $N^* < M^*$ (condition (A.i)). (We discharge this assumption in (b).) Consider the pure strategy $\sigma^*(q^*, N^*, M^*)$ with $q^* = 1.^{40}$ Since, in this section, all beliefs are derived according to Bayes' rule from the state-dependent stationary distributions of queue lengths induced by $\sigma^*(q^*, N^*, M^*)$, we use the shorthand $\bar{\mu}_n^t$ for $\bar{\mu}_n^t(q^*, N^*, M^*)$ to lighten notation.

Under $\sigma^*(1, N^*, M^*)$, so as to observe the first in line's behavior when the queue length reaches N^* , the *n*th in line must wait for $N^* - n + 1$ periods. Using (8) evaluated at $m = N^* - n + 1$ in (12), we have that the equilibrium payoff $U_n^*(\bar{\mu}_n^0)$ for an agent having joined the queue at the *n*th position in line and forming the belief $\bar{\mu}_n^0$ can be rewritten as

$$U_n^*(\bar{\mu}_n^0) = U_n(N^* - n + 1, \bar{\mu}_n^0) + \bar{\mu}_n^0(1 - \alpha)^{N^* - n + 1}\delta^{N^* - n + 1}(A_n - 1),$$

where, by (11), $A_n \ge 1$. The second term on the right-hand side reflects the informational benefit accruing to the *n*th agent in line after she has observed $N^* - n + 1$ failures, at which point she obtains the first in line's information.

A sufficient condition for the equilibrium condition (A.iv) is that for all $n = 2, ..., N^*$ and for all $m = 0, 1, ..., N^* - n + 1$,

$$U_n(N^* - n + 1, \bar{\mu}_n^0) \ge U_n(m, \bar{\mu}_n^0).$$

⁴⁰Our proof can easily be extended to mixed strategies with $q^* < 1$. Then a sufficient condition for (A.iv) is that for all $n = 2, ..., N^* + 1$ and for all $m = 0, 1, ..., N^* - n + 2$, $U_n(N^* - n + 2, \bar{\mu}_n^0) \ge U_n(m, \bar{\mu}_n^0)$ or, equivalently, that $\mathcal{N}(n, \bar{\mu}_n^0) \ge N^* - n + 2$ for every $n = 2, ..., N^* + 1$.

Theoretical Economics 14 (2019)

By Lemma 2, the above is equivalent to

$$\mathcal{N}(n, \bar{\mu}_n^0) \ge N^* - n + 1, \quad \forall n = 2, \dots, N^*.$$
 (A.31)

This sufficient condition ensures that the *n*th agent in line does not want to autonomously renege on the queue before the first in line has completed her N^* periods of experimentation, even if the informational benefit that accrues to her as she obtains the first in line's information is ignored. Equivalently, it says that based solely on her belief $\bar{\mu}_n^0$ formed upon joining the queue together with her own observation of the server (in)activity, the *n*th agent in line is willing to experiment for $N^* - n + 1$ periods.

We now provide a sufficient condition for (A.31). We proceed by induction for $1 \le n < N^* - 1$. We first establish a condition that ensures that if $\mathcal{N}(n, \bar{\mu}_n^0) \ge N^* - n + 1$, so that the *n*th in line does not want to autonomously renege on the queue before the first in line has completed her N^* periods of experimentation, then $\mathcal{N}(n + 1, \bar{\mu}_{n+1}^0) \ge N^* - n$ and the n + 1th in line also does not want to autonomously renege on the queue before the first the first in line has completed her N^* periods of experimentation.

Suppose that $\mathcal{N}(n, \bar{\mu}_n^0) \ge N^* - n + 1$. By Lemma 2, this is equivalent to⁴¹

$$\bar{\mu}_n^{N^*-n} \ge \underline{\mu}_n,\tag{A.32}$$

where $\underline{\mu}_n$ is the *n*th in line's cutoff belief defined in (10). Rewriting the above as likelihood ratios, using (5), (13), and (A.6) for the left-hand side, we obtain

$$\frac{\mu}{1-\mu}y_n N^* (1-\alpha)^{N^*-n} \ge \frac{\psi(1-\delta)}{\alpha(\psi^n \delta w - 1)}.$$
(A.33)

By the same token, the condition $\mathcal{N}(n+1, \bar{\mu}_{n+1}^0) \ge N^* - n$ is equivalent to

$$\bar{\mu}_{n+1}^{N^*-(n+1)} \ge \underline{\mu}_{n+1} \tag{A.34}$$

or, expressed as likelihood ratios,

$$\frac{\mu}{1-\mu}y_{n+1}N^*(1-\alpha)^{N^*-(n+1)} \ge \frac{\psi(1-\delta)}{\alpha(\psi^{n+1}\delta w - 1)}$$

Using the lower bound on y_{n+1} derived in (A.18) in the above inequality, we obtain a sufficient condition for (A.34):

$$\frac{\mu}{1-\mu}y_nN^*(1-\alpha)^{N^*-n} > \frac{\psi(1-\delta)\big(1-\alpha(1-\alpha)\big)}{\alpha(\psi^{n+1}\delta w - 1)}.$$

⁴¹ At the $(N^* - n + 1)$ th failure, the *n*th in line observes the first in line's behavior and updates her posterior to $\bar{\mu}_1^{N^*}$ if the first in line reneges or to 1 if the first in line does not renege. If $\bar{\mu}_n^{N^*-n+1} < \underline{\mu}_n$ then, based on her private learning alone, the *n*th in line would like to renege at the exit stage following the $(N^* - n + 1)$ th failure. However, because many rounds of exit are allowed at any exit stage, it is costless for the *n*th in line first to observe the first in line's behavior at the first round of the exit stage, as she is still able to renege at a later round of the same exit stage in case the first in line reneges, but benefits by staying in line in case the first in line does not renege.

Now compare this sufficient condition with (A.33). If we can show that

$$\frac{\psi(1-\delta)}{\alpha(\psi^n\delta w-1)} > \frac{\psi(1-\delta)(1-\alpha(1-\alpha))}{\alpha(\psi^{n+1}\delta w-1)},\tag{A.35}$$

then we have shown that (A.34) holds whenever (A.32) holds. Using $\psi = \alpha/(1 - \delta(1 - \alpha))$ and rearranging, (A.35) is equivalent to

$$\psi^n \delta w \frac{\alpha^2 - (1 - \delta)\psi}{\alpha^2} > 1$$

Depending on the sign of $\alpha^2/(\alpha^2 - (1 - \delta)\psi)$, we therefore have

(A.35)
$$\Leftrightarrow \begin{cases} \psi^n \delta w > \frac{\alpha^2}{\alpha^2 - (1 - \delta)\psi} & \text{if } \frac{\alpha^2}{\alpha^2 - (1 - \delta)\psi} > 0, \\ \psi^n \delta w < \frac{\alpha^2}{\alpha^2 - (1 - \delta)\psi} & \text{if } \frac{\alpha^2}{\alpha^2 - (1 - \delta)\psi} < 0. \end{cases}$$

Since $\psi^n \delta w > 0$ for all $(\alpha, \delta) \in (0, 1)^2$ and $1 \le n < N^*$, (A.35) holds if and only if

$$\psi^n \delta w > \frac{\alpha^2}{\alpha^2 - \psi(1 - \delta)} > 0. \tag{A.36}$$

The above gives a sufficient condition under which, if $\mathcal{N}(n, \bar{\mu}_n^0) = N^* - n + 1$, i.e., the *n*th in line does not find it optimal to renege before the first in line has completed her experimentation, then $\mathcal{N}(n+1, \bar{\mu}_{n+1}^0) \ge N^* - n$ and the n + 1th in line also does not want to renege before the first in line has completed her experimentation. This establishes the *n*th induction step.

We initiate the induction at n = 1 by observing that if $\mathcal{N}(1, \bar{\mu}_1^0) = N^*$, then the first in line's posterior beliefs after having observed N^* and $N^* + 1$ failures satisfy

$$\bar{\boldsymbol{\mu}}_1^{N^*-1} > \underline{\boldsymbol{\mu}}_1 \ge \bar{\boldsymbol{\mu}}_1^{N^*},$$

thereby satisfying (A.32), and initiating our induction.

Finally, observe that $\psi \delta w > \psi^2 \delta w > \cdots > \psi^{N^*-1} \delta w$. Thus, if (A.36) holds for $n = N^* - 1$, it holds for every step $1 \le n < N^* - 1$ of the induction.

In summary, we have shown that if $M^* = \mathcal{M}$ and $\mathcal{N}(1, \bar{\mu}_1^0(q^*, N^*, M^*)) = N^*$ (i.e., equilibrium conditions (A.ii) and (A.iii) hold), then a sufficient condition for (A.31) (and therefore for equilibrium condition (A.iv)) is that⁴²

$$\underbrace{\psi^{N^*} \delta w}_{V_{N^*}} > \underbrace{\frac{\alpha^2}{\alpha^2 - \psi(1 - \delta)}}_{f_2(\alpha, \delta)} > 0.$$
(A.37)

The above condition is given as (15) in Lemma 5.

⁴²Allowing for $q^* < 1$.

To finish, we highlight some properties of the function $f_2(\alpha, \delta)$ defined above. For every $(\alpha, \delta) \in (0, 1) \times (0, 1]$, $f_2(\alpha, \delta)$ is a strictly decreasing function of δ , with

$$\frac{\partial}{\partial \delta} f_2(\alpha, \delta) = -\left(\frac{\alpha \psi}{\alpha^2 - \psi(1 - \delta)}\right)^2 < 0.$$

Let

$$\delta_2(\alpha) := \frac{1 - \alpha}{1 - \alpha(1 - \alpha)} \tag{A.38}$$

be the value of δ setting $\alpha^2 - \psi(1 - \delta) = 0$. Then for every $\alpha \in (0, 1)$, $f_2(\alpha, \delta) > 0$ if and only if $\delta \in (\delta_2(\alpha), 1]$, with $\lim_{\delta \downarrow \delta_2(\alpha)} f_2(\alpha, \delta) = +\infty$ and $f_2(\alpha, 1) = 1$.

(b) From (4), we have that

$$\psi^{M^*} \delta w \ge 1 > \psi^{M^* + 1} \delta w.$$

Rewriting the second inequality, using $\psi > 0$, gives

$$\frac{1}{\psi} > \psi^{M^*} \delta w. \tag{A.39}$$

Using some algebra and $\psi = \alpha/(1 - \delta(1 - \alpha))$, it is easy to show that for every $\alpha \in (0, 1)$ and $\delta \in (\delta_2(\alpha), 1)$,

$$\frac{\alpha^2}{\alpha^2 - \psi(1 - \delta)} > \frac{1}{\psi}.$$

Using the above together with (A.39), we obtain that for every $\alpha \in (0, 1)$, if (A.37), then

$$\psi^{N^*} \delta w > \psi^{M^*} \delta w$$

or, equivalently, $N^* < M^*$.

A.6 Proof of Lemma 6

Let $M^* = \mathcal{M}$ and $N^* = \mathcal{N}(1, \bar{\mu}_0^1(q^*, N^*, M^*))$. We provide sufficient conditions on the primitives to ensure that (15) holds.

Consider the function $C : \mathbb{R}^+ \to \mathbb{R}^+$ defined by $C(x) := x(1-\alpha)^x$. It is quasiconcave in *x* and maximized at $x = -1/\ln(1-\alpha) > 0$. Evaluating the function at its maximum using $(1-\alpha)^{-1/\ln(1-\alpha)} = e^{-1}$, we have that for every $x \ge 0$,

$$C(x) \le \frac{-1}{e\ln(1-\alpha)}.\tag{A.40}$$

Now consider the function $\overline{C} : \mathbb{R}^+ \to \mathbb{R}^+$ defined by

$$\bar{C}(x) := \frac{-2}{e\ln(1-\alpha)} (1-\alpha)^{x/2}.$$
(A.41)

It is strictly decreasing in *x*. Moreover, it is evident from (A.40) and (A.41) that for every $x \ge 0$,

$$C(x) \le \bar{C}(x). \tag{A.42}$$

We now use the functions defined above to provide an equilibrium-independent upper bound on N^* . By Lemma 2, for $q^* \leq 1$, the equilibrium condition $N^* = \mathcal{N}(1, \bar{\mu}_1^0(q^*, N^*, M^*))$ is equivalent to

$$\bar{\mu}_1^{N^*}(q^*, N^*, M^*) \le \underline{\mu}_1 < \bar{\mu}_1^{N^*-1}(q^*, N^*, M^*),$$

where $\bar{\mu}_1^t(q^*, N^*, M^*)$ is the first in line's equilibrium posterior belief after *t* failures and $\underline{\mu}_1$ is the first in line's threshold belief defined in (10). Rewriting the second inequality above as likelihood ratios, using (5), (13), and (A.6) for the right-hand side, we obtain

$$\frac{\psi(1-\delta)}{\alpha(\psi\delta w-1)} < \frac{\mu}{1-\mu} y_1(q^*, N^*, M^*) (N^*+1-q^*)(1-\alpha)^{N^*-1}.$$

Rearranging gives

$$1 - \delta < y_1(q^*, N^*, M^*)(N^* + 1 - q^*)(1 - \alpha)^{N^* + 1} \frac{\mu}{1 - \mu} \frac{\alpha(\psi \delta w - 1)}{\psi(1 - \alpha)^2}$$

Since $y_1(q^*, N^*, M^*) \in [0, 1]$, $q^* \le 1$ and it follows that

$$1 - \delta < (N^* + 1)(1 - \alpha)^{N^* + 1} \frac{\mu}{1 - \mu} \frac{\alpha(\psi \delta w - 1)}{\psi(1 - \alpha)^2}.$$
 (A.43)

Finally, observe that since $\delta < 1$ and $\alpha < 1$,

$$\frac{\mu}{1-\mu}\frac{\alpha(\psi\delta w-1)}{\psi(1-\alpha)^2} < \frac{1}{(1-\alpha)^2}\frac{\alpha w\mu}{1-\mu}.$$

This, together with (A.43), implies

$$(1-\delta)(1-\alpha)^2 \frac{1-\mu}{\alpha w\mu} < C(N^*+1).$$
 (A.44)

By (16), the left-hand side above is strictly less than the upper bound on *C* given in (A.40). Consequently, the inequality above bounds $N^* + 1$, as illustrated in Figure 6, where $N^* + 1$ belongs to the highlighted interval. Choose \bar{N} so that

$$\bar{C}(\bar{N}) = (1-\delta)(1-\alpha)^2 \frac{1-\mu}{\alpha w\mu}$$
(A.45)

or, equivalently,

$$\bar{N} = 4 + \frac{\ln(1-\delta) - \ln\frac{w\alpha\mu}{1-\mu} + \ln\left[-\frac{e}{2}\ln(1-\alpha)\right]}{\frac{1}{2}\ln(1-\alpha)}.$$
(A.46)



FIGURE 6. The functions C(x) and $\overline{C}(x)$ for $\alpha = 0.4$.

Equations (A.44) and (A.45), together with (A.42) evaluated at $x = \bar{N}$ imply that $C(\bar{N}) \le C(N^* + 1)$. Since, by (16), \bar{N} lies on the decreasing portion of *C*, we have that $\bar{N} > N^* + 1$. Hence, $\bar{N} > N^*$.

We now show that the conditions of Lemma 6 are sufficient for (15). As shown at the end of the proof of Lemma 5(a), the second inequality in (15) is equivalent to $\delta > \delta_2(\alpha)$, where $\delta_2(\alpha)$ is defined in (A.38). Observe that $1 - \alpha^2 > \delta_2(\alpha)$. Therefore, the condition $\delta > 1 - \alpha^2$ (condition (17)) is sufficient for the second inequality in (15) to hold.

Now let us take logarithms of the first inequality in (15):

$$\ln(\delta w) + N^* \ln \psi + \ln \left(1 - \frac{\psi(1-\delta)}{\alpha^2} \right) > 0.$$
 (A.47)

Observe that since $\psi \in (0, 1)$, we have

$$0 \ge N^* \ln \psi \ge -N^* \frac{1-\psi}{\psi} = -N^*(1-\delta)\phi,$$

where the second inequality is an application of $\ln x \ge (x - 1)/x$, and where the final expression is obtained from (1) and the definition of ψ in (2). An application of $\ln x \ge (x - 1)/x$ also gives a lower bound on the final term in (A.47). Substitution of these bounds gives the following sufficient condition for the first inequality in (15):

$$\ln(\delta w) - N^*(1-\delta)\phi - \frac{1-\delta}{\delta(1-\alpha(1-\alpha)) - (1-\alpha)} > 0.$$

Substitute the upper bound \overline{N} defined in (A.46) for N^* and use (17) to get a lower bound on the denominator of the fraction above. This gives the new sufficient condition

$$\ln w > (1-\delta) \left(\bar{N}\phi + \frac{1}{\alpha^3(1-\alpha)} \right) - \ln \delta$$

or, substituting from (A.46) and simplifying,

$$\ln w > (1-\delta) \left(4\phi + \frac{\ln(1-\delta) - \ln\frac{w\alpha\mu}{1-\mu} + \ln\left[-\frac{e}{2}\ln(1-\alpha)\right]}{\frac{1}{2}\ln(1-\alpha)}\phi + \frac{1}{\alpha^3(1-\alpha)} \right)$$
$$-\ln\delta. \tag{A.48}$$

It is evident that the right-hand side is continuous and strictly decreasing in δ . It is positive for every $\delta \in (0, 1)$ and tends to zero when $\delta \rightarrow 1$. The left-hand side is a strictly positive constant, since w > 1. Therefore, by the intermediate value theorem, (A.48) imposes a lower bound on δ . In the lemma, this lower bound is labelled δ_0 .

A.7 Payoffs for the arrival at the M + 1th position under imperfect revelation

A.7.1 *Beliefs* Consider the profile $\sigma^*(q, N, M)$ with imperfect revelation, i.e., with M < N, $q \in (0, 1]$, or $M = N \ge 1$, $q \in (0, 1)$. An agent who arrives at the queue believing that it is in the steady state induced by $\sigma^*(q, N, M)$ bases her posterior beliefs about the server state and about whether the first in line is informed on the state-dependent invariant measures of queue lengths given in Appendix A.2.

When N > M + 1, $q \in (0, 1]$, or N = M + 1, $q \in (0, 1)$, the agent arriving at the queue at the M + 1th position does not know whether she is the first, second, third, etc. agent to arrive at that position behind the current first in line. If the current first in line joined the queue at the first position and subsequently observed M failures, the agent is the first arrival at the M + 1th position behind the current first in line. If the current first in line joined the queue at the first position and subsequently observed M + 1 failures, the agent is the second arrival at the M + 1th position behind the current first in line, and so on.

Generalizing, if the first in line observed *n* failures, n = M, ..., N, the agent is the n - M + 1th arrival at the M + 1th position. (If the first in line observes N + 1 failures (probability 1 - q), then she reneges at the N + 1th failure, so that there cannot be an N - M + 2th arrival at the M + 1th position.) Equivalently, if the first in line observed M - 1 + k failures, the agent is the *k*th arrival at the M + 1th position, k = 1, ..., N - M + 1. (Being the N - M + 1th arrival is possible only if q < 1 and the first in line's randomization makes her experiment for N + 1 periods, i.e., she does not renege at the *N*th service failure.)

Define $\overline{\xi}_k^{\text{good}}$ to be the probability that an agent who arrives at the queue at the M + 1th position attaches to being the *k*th arrival at the M + 1th position behind the current first in line and to that first in line being uninformed. The superscript indicates that this probability is conditioned on the server being good. We have

$$\bar{\xi}_{k}^{\text{good}} = \begin{cases} \frac{(1-\alpha)^{M+k-1}y_{1}}{y_{M+1}} & \text{for } k = 1, \dots, N-M, \\ (1-q)\frac{(1-\alpha)^{M+k-1}y_{1}}{y_{M+1}} & \text{for } k = N-M+1, \end{cases}$$
(A.49)

where $(1 - \alpha)^{M+k-1}y_1$ is the stationary probability that the current first in line joined the queue at the first position and subsequently observed M + k - 1 service failures, conditional on the server being good.

By the same token, conditional on the server being bad, we have

$$\bar{\xi}_{k}^{\text{bad}} = \begin{cases} \frac{x_{1}}{x_{M+1}} = \frac{1}{N - M + 1 - q} & \text{for } k = 1, \dots, N - M, \\ (1 - q)\frac{x_{1}}{x_{M+1}} = \frac{1 - q}{N - M + 1 - q} & \text{for } k = N - M + 1, \end{cases}$$
(A.50)

where x_1 is the stationary probability that the current first in line joined the queue at the first position (and subsequently observed M + k - 1 service failures; this event has probability 1 at a bad server), conditional on the server being bad.

When N = M + 1 and q = 1 or when M = N > 1 and $q \in (0, 1)$, the agent arriving at the M + 1th position knows that she is the first and only agent arriving at this position behind the current first in line. As in (A.49) and (A.50), she assigns probability $\bar{\xi}_1^{\text{good}} = (1 - q)(1 - \alpha)^M y_1 / y_{M+1}$ to the first in line being uninformed at a good server and probability $\bar{\xi}_1^{\text{bad}} = (1 - q)x_1 / x_{M+1} = 1$ to the first in line being uninformed at a bad server.

Thus, (A.49) and (A.50) describe the beliefs of the arrival at the M + 1th position regarding the first in line, both when N < M, $q \in (0, 1]$, and when M = N > 1, $q \in (0, 1)$.

Sections A.7.2 and A.7.3 describe the payoff to an agent arriving at the M + 1th position who deviates from the profile $\sigma^*(q, N, M)$ by joining the queue for m periods. To simplify the exposition, we assume that q = 1. In Section A.7.4 we derive the payoff for m = 1, allowing for q < 1, as this will be used in later sections.

A.7.2 Payoff from experimenting for $1 \le m \le N - M$ periods (for q = 1) We now describe the payoff, $\mathcal{U}_{M+1}(m, \bar{\mu}_{M+1}^0)$, of an agent who arrives at the queue as the M + 1th in line, forms the posterior belief $\bar{\mu}_{M+1}^0$ that the server is good as described in (13), and adopts the following strategy. Wait m periods. If a service event occurs during these m periods, remain in the queue until served. If before any service event occurs the first in line reneges on the queue, then renege in the same period as the first in line. In all other cases (i.e., there are m failures and the first in line has not reneged), autonomously renege after m periods. The details of $\mathcal{U}_{M+1}(m, \bar{\mu}_{M+1}^0)$ can be explained as follows.

Suppose the server is bad. Then the first in line will renege after having observed a total of *N* failures. If an agent is the *k*th arrival at the M + 1th position, she joins the queue when the first in line has already observed M - 1 + k failures. Thus, the first in line will stay in line for another N - (M - 1 + k) failures. The *k*th arrival's payoff depends on whether *m* is greater or less than this number.

If $m \ge N - M + 1 - k$ (or, equivalently, $k \ge N - M + 1 - m$), the agent who is the *k*th arrival at the M + 1th position reneges together with the first in line after having observed N - M + 1 - k failures and her payoff is $\delta^{N-M+1-k}$. If m < N - M + 1 - k (or, equivalently, k < N - M + 1 - m), she reneges autonomously after having observed *m* failures and her payoff is δ^m .

Therefore, the expected payoff at a bad server of an agent joining at the M + 1th position is

$$\sum_{k=1}^{N-M-m} \bar{\xi}_k^{\text{bad}} \delta^m + \sum_{k=N-M+1-m}^{N-M} \bar{\xi}_k^{\text{bad}} \delta^{N-M+1-k}.$$
(A.51)

Using the change of variable i = M + k gives

$$\sum_{i=M+1}^{N-m} \bar{\xi}_{i-M}^{\text{bad}} \delta^m + \sum_{i=N+1-m}^{N} \bar{\xi}_{i-M}^{\text{bad}} \delta^{N+1-i}.$$
(A.52)

Now suppose that the server is good. If the first in line is uninformed, then she will renege after having observed a total of N failures. If an agent is the kth arrival at the M + 1th position, she joins the queue when the first in line has already observed M – 1 + k failures. Thus, the uninformed first in line will stay in line for another N - (M - 1 + k) failures. The kth arrival's payoff depends on whether m is greater or less than this number.

If $m \ge N - M + 1 - k$ (or, equivalently, $k \ge N - M + 1 - m$), the agent who is the *k*th arrival at the M + 1th position reneges together with the first in line after having observed N - M + 1 - k failures unless service occurs first. Thus, her payoff is

$$\sum_{s=1}^{N-M+1-k} \alpha (1-\alpha)^{s-1} \delta^s W_{M+1} + (1-\alpha)^{N-M+1-k} \delta^{N-M+1-k}.$$
 (A.53)

If m < N - M + 1 - k (or, equivalently, k < N - M + 1 - m), the *k*th arrival at the M + 1th position reneges autonomously after having observed *m* failures unless service occurs first. Her payoff is

$$\sum_{s=1}^{m} \alpha (1-\alpha)^{s-1} \delta^s W_{M+1} + (1-\alpha)^m \delta^m.$$
 (A.54)

If the first in line is informed, she never reneges on the queue, and the agent who is the *k*th arrival at the M + 1th position reneges autonomously after having observed m failures, so her payoff is (A.54). Therefore, the expected payoff at a good server of an agent joining at the M + 1th position is

$$\begin{split} &\sum_{k=1}^{N-M-m} \bar{\xi}_{k}^{\text{good}} \left(\sum_{s=1}^{m} \alpha (1-\alpha)^{s-1} \delta^{s} W_{M+1} + (1-\alpha)^{m} \delta^{m} \right) \\ &+ \sum_{k=N-M+1-m}^{N-M} \bar{\xi}_{k}^{\text{good}} \left(\sum_{s=1}^{N-M+1-k} \alpha (1-\alpha)^{s-1} \delta^{s} W_{M+1} + (1-\alpha)^{N-M+1-k} \delta^{N-M+1-k} \right) \\ &+ \left(1 - \sum_{k=1}^{N-M} \bar{\xi}_{k}^{\text{good}} \right) \left(\sum_{s=1}^{m} \alpha (1-\alpha)^{s-1} \delta^{s} W_{M+1} + (1-\alpha)^{m} \delta^{m} \right), \end{split}$$

where $1 - \sum_{k=1}^{N-M} \bar{\xi}_k^{\text{good}}$ is the probability that the arrival at the M + 1th position attributes to the first in line being informed at a good server. Factorizing the term in the large parentheses in the first line and adding up with the third line gives

$$\sum_{k=N-M+1-m}^{N-M} \bar{\xi}_{k}^{\text{good}} \left(\sum_{s=1}^{N-M+1-k} \alpha (1-\alpha)^{s-1} \delta^{s} W_{M+1} + (1-\alpha)^{N-M+1-k} \delta^{N-M+1-k} \right) \\ + \left(1 - \sum_{k=N-M+1-m}^{N-M} \bar{\xi}_{k}^{\text{good}} \right) \left(\sum_{s=1}^{m} \alpha (1-\alpha)^{s-1} \delta^{s} W_{M+1} + (1-\alpha)^{m} \delta^{m} \right).$$

Using the change of variable i = M + k gives

$$\sum_{i=N+1-m}^{N} \bar{\xi}_{i-M}^{\text{good}} \left(\sum_{s=1}^{N+1-i} \alpha (1-\alpha)^{s-1} \delta^{s} W_{M+1} + (1-\alpha)^{N+1-i} \delta^{N+1-i} \right) + \left(1 - \sum_{i=N+1-m}^{N} \bar{\xi}_{i-M}^{\text{good}} \right) \left(\sum_{s=1}^{m} \alpha (1-\alpha)^{s-1} \delta^{s} W_{M+1} + (1-\alpha)^{m} \delta^{m} \right).$$
(A.55)

From (A.52) and (A.55) we have that

$$\begin{split} \mathcal{U}_{M+1}(m, \bar{\mu}_{M+1}^{0}) \\ &= \left(1 - \bar{\mu}_{M+1}^{0}\right) \left\{ \sum_{i=M+1}^{N-m} \bar{\xi}_{i-M}^{\text{bad}} \delta^{m} + \sum_{i=N+1-m}^{N} \bar{\xi}_{i-M}^{\text{bad}} \delta^{N+1-i} \right\} \\ &+ \bar{\mu}_{M+1}^{0} \left\{ \sum_{i=N+1-m}^{N} \bar{\xi}_{i-M}^{\text{good}} \left(\sum_{s=1}^{N+1-i} \alpha (1-\alpha)^{s-1} \delta^{s} W_{M+1} + (1-\alpha)^{N+1-i} \delta^{N+1-i} \right) \\ &+ \left(1 - \sum_{i=N+1-m}^{N} \bar{\xi}_{i-M}^{\text{good}} \right) \left(\sum_{s=1}^{m} \alpha (1-\alpha)^{s-1} \delta^{s} W_{M+1} + (1-\alpha)^{m} \delta^{m} \right) \right\}. \end{split}$$

Finally, using (A.49) and (A.50) gives

$$\begin{aligned} \mathcal{U}_{M+1}(m, \bar{\mu}_{M+1}^{0}) \\ &= \frac{1 - \bar{\mu}_{M+1}^{0}}{N - M} \left\{ \sum_{i=M+1}^{N-m} \delta^{m} + \sum_{i=N+1-m}^{N} \delta^{N+1-i} \right\} \\ &+ \bar{\mu}_{M+1}^{0} \left\{ \sum_{i=N+1-m}^{N} \frac{(1 - \alpha)^{i-1} y_{1}}{y_{M+1}} \left(\sum_{s=1}^{N+1-i} \alpha (1 - \alpha)^{s-1} \delta^{s} W_{M+1} + (1 - \alpha)^{N+1-i} \delta^{N+1-i} \right) \right. \\ &+ \left(1 - \sum_{i=N+1-m}^{N} \frac{(1 - \alpha)^{i-1} y_{1}}{y_{M+1}} \right) \left(\sum_{s=1}^{m} \alpha (1 - \alpha)^{s-1} \delta^{s} W_{M+1} + (1 - \alpha)^{m} \delta^{m} \right) \right\}. \end{aligned}$$
(A.56)

A.7.3 Payoff from joining the queue for m = N - M periods and herding on the first in line (for q = 1) We describe the payoff, $U_{M+1}^h(\bar{\mu}_{M+1}^0)$, of an agent who arrives at the queue as the M + 1th in line, forms the posterior belief $\bar{\mu}_{M+1}^0$ that the server is good as described in (13), and adopts the following strategy. Wait m = N - M periods. If a service event occurs during these m periods, remain in the queue until served. If before any service event occurs the first in line reneges on the queue, then renege in the same period as the first in line. In all other cases (i.e., there are m failures and the first in line has not reneged), remain in the queue until served.

Now, if the first in line does not renege after the *m*th failure, the agent at the M + 1th position learns that she was the first arrival at the M + 1th position behind the current first in line. She also learns that the first in line has not reneged despite having observed a total of N service failures. This means that the first in line must have previously observed service (i.e., she is informed). The agent at the M + 1th position therefore learns that the server is good and remains in line until served. The details of $U_{M+1}^h(\tilde{\mu}_{M+1}^0)$ can be explained as follows.

Suppose the server is bad. Then the possible contingencies faced by the *k*th arrival at the M + 1th position are as in Section A.7.2. Therefore, using m = N - M in (A.51), we obtain that the expected payoff at a bad server of an agent joining at the M + 1th position is

$$\sum_{k=1}^{N-M} \bar{\xi}_k^{\text{bad}} \delta^{N-M+1-k} = \sum_{i=M+1}^N \bar{\xi}_{i-M}^{\text{bad}} \delta^{N+1-i},$$
(A.57)

where the second expression is obtained using the change of variable i = M + k.

Now suppose that the server is good. If the first in line is uninformed, then she will renege after observing N failures. Since m = N - M, for every $k \ge 1$, the kth arrival at the M + 1th position reneges together with the first in line, so her payoff is (A.53) evaluated at m = N - M. If the first in line is informed, she never reneges on the queue. After observing the first in line not reneging despite N - M failures, the kth arrival at the M + 1th position learns that the server is good and remains in line until served. Hence, her payoff is

$$\sum_{s=1}^{N-M} \alpha (1-\alpha)^{s-1} \delta^s W_{M+1} + (1-\alpha)^{N-M} \delta^{N-M} V_{M+1} = V_{M+1},$$

where the equality follows from (2) and (6). Therefore, the expected payoff at a good server of an agent joining at the M + 1th position is

$$\sum_{k=1}^{N-M} \bar{\xi}_{k}^{\text{good}} \left(\sum_{s=1}^{N-M+1-k} \alpha (1-\alpha)^{s-1} \delta^{s} W_{M+1} + (1-\alpha)^{N-M+1-k} \delta^{N-M+1-k} \right) + \left(1 - \sum_{k=1}^{N-M} \bar{\xi}_{k}^{\text{good}} \right) V_{M+1}.$$

Using the change of variable i = M + k gives

$$\sum_{i=M+1}^{N} \bar{\xi}_{i-M}^{\text{good}} \left(\sum_{s=1}^{N+1-i} \alpha (1-\alpha)^{s-1} \delta^{s} W_{M+1} + (1-\alpha)^{N+1-i} \delta^{N+1-i} \right) + \left(1 - \sum_{i=M+1}^{N} \bar{\xi}_{i-M}^{\text{good}} \right) V_{M+1}.$$
(A.58)

Combining (A.57) and (A.58), we have that

$$\begin{aligned} \mathcal{U}_{M+1}^{h}(\bar{\mu}_{M+1}^{0}) &= \left(1 - \bar{\mu}_{M+1}^{0}\right) \sum_{i=M+1}^{N} \bar{\xi}_{i-M}^{\text{bad}} \delta^{N+1-i} \\ &+ \bar{\mu}_{M+1}^{0} \left\{ \sum_{i=M+1}^{N} \bar{\xi}_{i-M}^{\text{good}} \left(\sum_{s=1}^{N+1-i} \alpha (1 - \alpha)^{s-1} \delta^{s} W_{M+1} + (1 - \alpha)^{N+1-i} \delta^{N+1-i} \right) \\ &+ \left(1 - \sum_{i=M+1}^{N} \bar{\xi}_{i-M}^{\text{good}} \right) V_{M+1} \right\}. \end{aligned}$$

Finally, using (A.49) and (A.50) gives

$$\begin{aligned} \mathcal{U}_{M+1}^{h}(\bar{\mu}_{M+1}^{0}) \\ &= \frac{1 - \bar{\mu}_{M+1}^{0}}{N - M} \sum_{i=M+1}^{N} \delta^{N+1-i} \\ &+ \bar{\mu}_{M+1}^{0} \Biggl\{ \sum_{i=M+1}^{N} \frac{(1 - \alpha)^{i-1} y_{1}}{y_{M+1}} \Biggl(\sum_{s=1}^{N+1-i} \alpha (1 - \alpha)^{s-1} \delta^{s} W_{M+1} + (1 - \alpha)^{N+1-i} \delta^{N+1-i} \Biggr) \\ &+ \Biggl(1 - \sum_{i=M+1}^{N} \frac{(1 - \alpha)^{i-1} y_{1}}{y_{M+1}} \Biggr) V_{M+1} \Biggr\}. \end{aligned}$$
(A.59)

A.7.4 Payoff from experimenting for m = 1 period (for q < 1) Consider the deviation where the agent arriving at the $M^* + 1$ th position joins the queue for m = 1 period and reneges if the one service event she witnesses produces a failure, regardless of the first in line's behavior.

Suppose the server is bad. Then the M + 1th in line reneges after observing one failure, regardless of whether the first in line reneges in the same period as her. Consequently, the belief of the M + 1th in line regarding the first in line is irrelevant and her payoff is δ .

Now suppose that the server is good. The M + 1th in line reneges after observing one failure, regardless of whether the first in line reneges in the same period as her, unless service occurs first, in which case she remains in line until served. Consequently, the belief of the M + 1th in line regarding the first in line is irrelevant and her payoff is $\delta[\alpha W_{M+1} + 1 - \alpha]$.

Therefore,

$$\mathcal{U}_{M+1}(1,\bar{\mu}^0_{M+1}) = (1-\bar{\mu}^0_{M+1})\delta + \bar{\mu}^0_{M+1}\delta[\alpha W_{M+1} + 1 - \alpha]$$
(A.60)

$$= U_{M+1} (1, \bar{\mu}_{M+1}^0), \tag{A.61}$$

where the second equality follows from the definition of $U_{M+1}(1, \bar{\mu}_{M+1}^0)$ in (8).

A.8 Proof of Lemma 7

When $N^* = M^* = 1$, (A.4) and (A.6) give

$$y_1 = \frac{1}{1 + (1 - q^*)(1 - \alpha)}, \qquad y_2 = 1 - y_1; \qquad x_1 = \frac{1}{2 - q^*}, \qquad x_2 = 1 - x_1.$$

Using (5) and (13) gives

$$\begin{split} &\frac{\bar{\mu}_1^0}{1-\bar{\mu}_1^0} = \frac{\mu}{1-\mu} \frac{y_1}{x_1} = \frac{\mu}{1-\mu} \frac{2-q}{1+(1-q)(1-\alpha)},\\ &\frac{\bar{\mu}_2^0}{1-\bar{\mu}_2^0} = \frac{\mu}{1-\mu} \frac{1-y_1}{1-x_1} = \frac{\mu}{1-\mu} \frac{2-q}{1+(1-q)(1-\alpha)}(1-\alpha),\\ &\frac{\bar{\mu}_1^1}{1-\bar{\mu}_1^1} = \frac{\bar{\mu}_1^0}{1-\bar{\mu}_1^0}(1-\alpha), \end{split}$$

so that

$$\bar{\mu}_1^0 > \bar{\mu}_1^1 = \bar{\mu}_2^0 > 0 \tag{A.62}$$

for all $(\alpha, \mu, q) \in (0, 1)^3$.

Consider first the equilibrium condition (B.ii). For $M^* = N^* = 1$, this yields two conditions. (Since the first in line is necessarily uninformed, there is no third condition.) First, $U_1^*(\bar{\mu}_1^0) \ge 1$, where $U_1^*(\bar{\mu}_1^0)$, defined in (12), is the first in line's payoff from adhering to the strategy σ^* . But this is just the payoff from experimenting for one period. Indeed, since $A_1 = 1$ (by (11)) we have that $U_1^*(\bar{\mu}_1^0) = U_1(1, \bar{\mu}_1^0)$, where $U_1(1, \bar{\mu}_1^0)$ is defined in (8). Second, $\mathcal{U}_2(1, \bar{\mu}_2^0) \le 1$, where $\mathcal{U}_2(1, \bar{\mu}_2^0)$ is the second in line's payoff from deviating from the equilibrium strategy and experimenting for one period. By (A.61), $\mathcal{U}_2(1, \bar{\mu}_2^0) = U_2(1, \bar{\mu}_2^0)$. From Lemma 2 we have that $U_1(1, \bar{\mu}_1^0) \ge 1$ and $U_2(1, \bar{\mu}_2^0) \le 1$ if and only if $\bar{\mu}_1^0 \ge \underline{\mu}_1$ and $\bar{\mu}_2^0 \le \underline{\mu}_2$. Consider now the equilibrium condition (B.iii). From Lemma 2 we have that $q^* \in \underline{\mu}_1$.

Consider now the equilibrium condition (B.iii). From Lemma 2 we have that $q^* \in (0, 1)$ and $N^* = 1$ if and only if $\bar{\mu}_1^0 \ge \underline{\mu}_1 = \bar{\mu}_1^1$. Similarly, $q^* = 1$ and $N^* = 1$ if and only if $\bar{\mu}_1^0 = \mu \ge \underline{\mu}_1 > \bar{\mu}_1^1 = \mu(1 - \alpha)/(1 - \alpha\mu)$.

Finally, note that condition (B.iv) does not apply as no agent joins the queue at positions 2, 3, . . . in equilibrium, since $M^* + 1 = 2$.

In summary, the following conditions are necessary and sufficient for an equilibrium with $M^* = N^* = 1$ and $q^* \in (0, 1]$:

$$\bar{\mu}_1^0 \ge \underline{\mu}_1 \ge \bar{\mu}_1^1, \qquad \underline{\mu}_2 \ge \bar{\mu}_2^0. \tag{A.63}$$

For the existence of an equilibrium with $M^* = N^* = 1$ and $q^* \in (0, 1)$, the condition $\bar{\mu}_1^1 = \underline{\mu}_1$ is necessary and sufficient. (By (10), $\underline{\mu}_2 > \underline{\mu}_1$. Thus, (A.62) ensures that the other inequalities in (A.63) hold.) For each $(\alpha, \mu, q) \in (0, 1)^3$, the equation $\bar{\mu}_1^1 = \underline{\mu}_1$ has the unique solution $\delta_1(\alpha, \mu, q)$, which is defined in (20) and takes values in the interval $D_2^*(\alpha, \mu) \subset (0, 1)$, defined in (21).

For the existence of an equilibrium with $M^* = N^* = 1$ and $q^* = 1$, the condition $\bar{\mu}_1^0 \ge \underline{\mu}_1 > \bar{\mu}_1^1$ is necessary and sufficient. At each (α, μ) , this last pair of inequalities holds if and only if $\delta \in D_1^*(\alpha, \mu)$, defined in (19).

A.9 Proof of Lemma 8

By Lemma 7, the strategy $\sigma^*(q^*, 1, 1)$ cannot constitute an equilibrium if $\delta \ge 1/(1 + c)$, where c > 0 is defined in (21). This establishes the result for $M^* = N^* = 1$.

Now consider the strategy $\sigma^*(q^*, N^*, M^*)$, where either $M^* < N^*$ and $q^* \in (0, 1]$ or $M^* = N^* > 1$ and $q^* \in (0, 1)$. Suppose, by way of contradiction, that for every $\bar{\delta} \in (0, 1)$, there exists $\delta \in (\bar{\delta}, 1)$ at which $\sigma^*(q^*, N^*, M^*)$ constitutes an equilibrium with imperfect revelation. The agent arriving at the $M^* + 1$ th position cannot learn, merely by observing the queue length, that the server is in the good state. The equilibrium value M^* therefore satisfies the conditions given in (B.ii).

We show that this condition is violated as δ becomes arbitrarily close to 1 by showing that there exists a profitable deviation. The argument is that as δ approaches unity, joining the queue for one period becomes virtually costless and might ultimately result in service. Balking can therefore never be optimal.

Consider the deviation where the agent arriving at the M^* + 1th position joins the queue for m = 1 period and reneges if the one service event she witnesses produces a failure, regardless of the first in line's behavior. The payoff from this deviation for the agent arriving at the M^* + 1th position is $\mathcal{U}_{M^*+1}(1, \bar{\mu}^0_{M^*+1}(q^*, N^*, M^*))$, which is defined in (A.60).

We now show that when δ is arbitrarily close to 1, the agent arriving at the M^* + 1th position gets a lower payoff from balking than from joining the queue for one period and reneging if the one service event she witnesses produces a failure. From (6), we have that $\lim_{\delta \to 1} W_{M^*+1} = w$. Therefore, evaluating the expected payoff in (A.60) as δ becomes arbitrarily close to 1 gives

$$\lim_{\delta \to 1} \mathcal{U}_{M^*+1}(1, \bar{\mu}^0_{M^*+1}) = 1 - \bar{\mu}^0_{M^*+1} + \bar{\mu}^0_{M^*+1}[\alpha w + 1 - \alpha].$$

Using w > 1 in the above expression, it is immediate that, for every $\bar{\mu}_{M^*+1}^0 \in (0, 1)$,

$$\lim_{\delta \to 1} \mathcal{U}_{M^*+1}(1, \bar{\mu}^0_{M^*+1}) > 1,$$

violating the second condition in (B.ii). From (A.60), it is evident that $\mathcal{U}_{M^*+1}(1, \bar{\mu}_{M^*+1}^0)$ is continuous in δ . Hence, for every (α, μ) there exists a threshold $d(\alpha, \mu) < 1$ such that deviating from $\sigma^*(q^*, N^*, M^*)$ is profitable at every $\delta > d(\alpha, \mu)$, a contradiction.

A.10 Proof of Lemma 9

Assume that M > N > 1 and $q = 1.^{43}$ Let us derive a condition for $\bar{\mu}_1^0 < \mu$. From (13) we have that $\bar{\mu}_1^0 < \mu$ if and only if $Ny_1 < 1$. Dividing both sides of (A.19) by N gives

$$\frac{1}{Ny_1} = 1 + \frac{\frac{\phi^N - \phi^{M+1}}{1 - \phi} \left(1 - \frac{1 - \phi^{N+1}}{(1 - \phi)(1 + \phi)^N}\right) - \sum_{i=0}^{N-1} \frac{1 - \phi^i}{1 - \phi} \left(1 - \phi + \phi \left(\frac{\phi}{1 + \phi}\right)^N\right)}{N}.$$

A necessary and sufficient condition for $Ny_1 < 1$ is that the fraction on the right-hand side be strictly positive. Recall that $\phi > 0$. Since M > N, the first fraction at the numerator is positive. Since $(1 + \phi)^N > \sum_{i=0}^N \phi^i = (1 - \phi^{N+1})/(1 - \phi)$, the first term in parentheses at the numerator is also positive. Hence, the first summand in the numerator is positive. Therefore, a necessary and sufficient condition for $Ny_1 < 1$ is that

$$\left(\frac{\sum_{i=0}^{N-1} \frac{1-\phi^{i}}{1-\phi}}{\frac{\phi^{N}-\phi^{M+1}}{1-\phi}}\right) \frac{1-\phi+\frac{\phi^{N+1}}{(1+\phi)^{N}}}{1-\frac{1-\phi^{N+1}}{(1-\phi)(1+\phi)^{N}}} < 1.$$
(A.64)

The left-hand side of (A.64) is negative if $1 - \phi + \phi (\phi/(1+\phi))^N < 0$, since $\sum_{i=0}^{N-1} (1-\phi^i)/(1-\phi) > 0$ and we have already argued that the remaining two terms are positive. But $1 - \phi + \phi (\phi/(1+\phi))^N$ is decreasing in N and it is negative for N = 2 if and only if $1 + \phi - \phi^2 < 0$. This holds if and only if $\phi > (1 + \sqrt{5})/2$ or, by (1), $\alpha < 2/(3 + \sqrt{5})$. Hence, $\alpha < 2/(3 + \sqrt{5})$ is a sufficient condition for $Ny_1 < 1$ and $\overline{\mu}_1^0 < \mu$ when M > N > 1.

We now derive a condition for $\bar{\mu}_1^0 > \mu$ (or $Ny_1 > 1$) when M > N > 1 and q = 1.

Assume that $\phi < 1$ so that $a(N) := (1 - \phi)(1 + \phi)^N + \phi^{N+1}$ is positive. The unbracketed ratio in (A.64) equals $(1 - \phi)a(N)/(a(N) - 1)$. From $a(N + 1) - a(N) = (\phi - 1)\phi((1 + \phi)^N - \phi^N) > 0$ we have that a(N) is a strictly increasing function of N, so that a(N) > a(1) = 1. Therefore, a(N)/(a(N) - 1) > 1 and $(1 - \phi)a(N)/(a(N) - 1) > 1 - \phi$. Substitution of this lower bound then gives a sufficient condition for the left-hand side of (A.64) to exceed unity:

$$(1-\phi)\sum_{i=0}^{N-1} (1-\phi^i) > \phi^N - \phi^{M+1}$$
 or $(1-\phi)N > 1-\phi^M$.

The right-hand side is less than 1 for every M > 1. For $\phi < 1/2$, the left-hand side is greater than 1 for every N > 1. Thus, $\phi < 1/2$ or $\alpha > 2/3$ is sufficient for this condition. In summary, when $\alpha > 2/3$ and M > N > 1, the left-hand side of (A.64) exceeds unity, and $\bar{\mu}_1^0 > \mu$.

⁴³The proof of Lemma 9 also holds for the case 0 < q < 1. This is because the inequality (A.64) can be written as a linear function of q. Therefore, ensuring it holds for q = 1 and q = 0 (the latter is equivalent to setting q = 1 and increasing N by 1), which is what our proof does, is enough to ensure it holds for all q. The same argument holds for the case $M \le N$.

Finally, we derive a condition for $\bar{\mu}_1^0 > \mu$ (or $Ny_1 > 1$) when $M \le N$ and q = 1. When M < N and q = 1, substitutions from (A.4) and (A.5) give

$$\frac{1}{(M+1)y_1} = \frac{\frac{1-\phi^{M+1}}{1-\phi} - k_N(M+\phi^{-1})}{(M+1)(1-k_N)}$$
$$= 1 + \frac{\sum_{i=0}^{M} (\phi^i - 1) + k_N(1-\phi^{-1})}{(M+1)(1-k_N)}.$$

If $\phi < 1$ (and hence $0 < k_N < 1$) the numerator of the final fraction is negative, which implies that $y_1 > 1/(M + 1)$. Since $M + 1 \le N$, this implies that $y_1 > 1/N$ and that the sufficient condition for $\bar{\mu}_1^0 > \mu$ holds if $\phi < 1$. As $\alpha > 2/3$ is sufficient for $\phi < 1$, it is also sufficient for $\bar{\mu}_1^0 > \mu$ if M < N.

When M = N and q = 1, substitutions from (A.4) and (A.5) give

$$\frac{1}{Ny_1} = \frac{\frac{1-\phi^{N+1}}{1-\phi} - k_N(N+\phi^{-1})}{N(1-k_N)}$$

For $\phi < 1$, the condition $Ny_1 > 1$ can be written as

$$\frac{k_N}{\phi} > \frac{1-\phi^{N+1}}{1-\phi} - N.$$

A sufficient condition for this is $0 > (1 - \phi^{N+1})/(1 - \phi) - N$ for all N > 1. The right-hand side equals $1 + \sum_{i=0}^{N} (\phi^{i} - 1)$ and is strictly decreasing in N. Therefore, the condition is tightest for N = 2 and a sufficient condition is $1 > \phi + \phi^{2}$. This holds if $\phi > (\sqrt{5} - 1)/2$ or, by (1), if $\alpha > 2/(1 + \sqrt{5})$, which is ensured by the condition $\alpha > 2/3$.

In summary, when $\alpha > 2/3$ and $M \le N$, $\bar{\mu}_1^0 > \mu$.

A.11 Proof of Lemma 10

Let $M^* = \mathcal{M}$ and assume that $q^* = 1$. We are interested in the behavior as $\delta \to 1$ of equation (A.iii), which determines the equilibrium value N^* . Looking at the continuous version of this (i.e., treating both N^* and M^* as elements of \mathbb{R}^+), we have, by Lemma 2, that $\bar{\mu}_1^{N^*}(1, N^*, M^*) = \underline{\mu}_1$. Rewriting this equality as likelihood ratios, and using (5) and (13) gives

$$\frac{\mu}{1-\mu} y_1(1, N^*, M^*) N^* (1-\alpha)^{N^*} = \frac{\psi(1-\delta)}{\alpha(\psi\delta w - 1)}.$$
(A.65)

By Lemma 8, when $\delta \to 1$, we necessarily have $M^* > N^*$. We therefore use the expression for $y_1(1, N^*, M^*)$ calculated in (A.21). Substituting this and $\phi = \frac{1-\alpha}{\alpha}$ on the left, the above

equation becomes

$$\frac{N^* \left(\frac{\phi}{1+\phi}\right)^{N^*}}{\frac{1-\phi^{M^*+1}}{1-\phi} - \left(\frac{\phi}{1+\phi}\right)^{N^*} \left[1+\phi\sum_{i=0}^{N^*}\frac{1-\phi^{M^*-i}}{1-\phi}\right]} = \frac{\psi(1-\delta)}{\alpha(\psi\delta w-1)}\frac{1-\mu}{\mu}.$$

Finally, a lengthy rearranging gives

$$\theta_1(N^*)(\phi - 1) = \frac{\psi(1 - \delta)}{\alpha(\psi \delta w - 1)} \frac{1 - \mu}{\mu} \big[\phi^{M^* + 1} \kappa \big(N^* \big) - \theta \big(N^* \big) \big], \tag{A.66}$$

where $\kappa(N) := 1 - \frac{1}{(1+\phi)^N} \frac{1-\phi^{N+1}}{1-\phi}$, $\theta_1(N) := N(\frac{\phi}{1+\phi})^N$, and $\theta(N) := 1 - (\frac{\phi}{1+\phi})^N - \phi \theta_1(N)$. These functions satisfy, for all $N \ge 0$, the inequalities

$$\kappa(N) \in [0,1], \qquad \theta_1(N) \le \left[\ln\left(\frac{1+\phi}{\phi}\right) \right]^{-1}, \qquad 1 \ge \theta(N) \ge -\phi \left[\ln\left(\frac{1+\phi}{\phi}\right) \right]^{-1}.$$

(The last two bounds apply the inequality $N\beta^N \leq -1/\ln\beta$, which holds for any $\beta < 1$ and N > 0.)

We study the values of N^* that solve (A.66) when M^* is defined by $\psi^{M^*} \delta w = 1$ (i.e., the continuous version of condition (A.ii)), as $\delta \to 1$. By Lemma 4, a solution to this equation exists. We will use the fact that $M^* \to \infty$ as $\delta \to 1$. This follows as

$$M^* = \frac{\ln \delta w}{-\ln \psi} \approx \frac{\ln w}{1 - \psi} = \frac{\ln w}{(1 - \delta)\phi\psi} \approx \frac{\ln w}{(1 - \delta)\phi}.$$
 (A.67)

These approximations become arbitrarily good as $\delta \rightarrow 1$.

(a) Let $1/2 < \alpha < 1$ or, equivalently, let $0 < \phi < 1$. By (A.67), $M^* \to \infty$ as $\delta \to 1$. If $\phi < 1$, $\delta \to 1$, and $M^* \to \infty$, the right-hand side of (A.66) converges uniformly to zero for all values of $N^* \ge 1$. This implies that at any solution to (A.66), the left-hand side of (A.66) also converges to zero, so $N^* \to \infty$.

(b) Let $0 < \alpha < 1/2$ or, equivalently, $\phi > 1$. We begin by showing that the right-hand side of (A.66) tends to infinity as $\delta \rightarrow 1$ for all $N^* > 1$. First observe that $\kappa(N^*) > 0$ when $N^* > 1$. Now if we substitute from (A.67) for M^* , the right-hand side of (A.66) becomes

$$K\frac{\psi}{\psi\delta w-1}\Big[(1-\delta)\phi^{\frac{\ln w}{(1-\delta)\phi}+1}\kappa(N^*)-(1-\delta)\theta(N^*)\Big],$$

where *K* is independent of δ and *N*. Since $\phi > 1$ and w > 1, this tends to infinity for all $N^* > 1$ as $\delta \to 1$. This is because the term that multiplies the square brackets tends to 1/(w-1) and the first term in the square brackets tends to infinity. However, observe that the left-hand side of (A.66) is bounded above, since $\theta_1(N^*)$ has an upper bound, so the equality cannot hold at $N^* > 1$ when $\delta \to 1$.

However, when $N^* = 1$, $\kappa(N^*) = 0$, so the right-hand side of (A.66) converges to zero as $\delta \to 1$ when $N^* = 1$. This, combined with the continuity of (A.66) in N^* , implies that N^* that solves (A.66) converges to unity as $\delta \to 1$.

(c) For $\alpha = 1/2$ (or $\phi = 1$), the stationary distribution of queue lengths is given in Lemma B.1 in the Supplemental Material. Using q = 1 and factorizing M, we have

$$\frac{1}{y_1(1,N,M)} = M\left(1 - (N+1)2^{-N}\right) + 1 - 2^{-N}\left(1 - \frac{N(N+1)}{2}\right).$$

Using this in (A.65), evaluated at $\alpha = 1/2$, gives

$$N^{*}2^{-N^{*}} = \frac{2\psi(1-\delta)}{\psi\delta w - 1} \frac{1-\mu}{\mu} \bigg[M^{*} \big(1 - (N^{*}+1)2^{-N^{*}}\big) + 1 - 2^{-N^{*}} \bigg(1 - \frac{N^{*}(N^{*}+1)}{2}\bigg) \bigg].$$

Making the substitution (A.67) for M^* and letting $\delta \to 1$ (observing that $(2\psi \ln w)/[\phi(\psi \delta w - 1)] \to (2\ln w)/(w - 1)$ and that all relevant terms are bounded) gives the relation for the limiting values of N^* :

$$\frac{N^*}{2^{N^*} - N^* - 1} = \frac{2\ln w}{w - 1} \frac{1 - \mu}{\mu}.$$

A finite solution, $c(\mu)$, to this exists, as the left-hand side above is continuous and strictly decreasing for $N^* > 1$, is unbounded for $N^* = 1$, and converges to zero as $N^* \to \infty$. The solution is also, therefore, strictly larger than unity.

References

Altman, Eitan and Nahum Shimkin (1998), "Individual equilibrium and learning in processor sharing systems." *Operations Research*, 46, 776–784. [651]

Assaf, David and Moshe Haviv (1990), "Reneging from processor sharing systems and random queues." *Mathematics of Operations Research*, 15, 129–138. [651]

Banerjee, Abhijit and Drew Fudenberg (2004), "Word-of-mouth learning." *Games and Economic Behavior*, 46, 1–22. [650]

Banerjee, Abhijit V. (1992), "A simple model of herd behavior." *Quarterly Journal of Economics*, 107, 797–817. [650]

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), "A theory of fads, fashion, custom, and cultural change as informational cascades." *Journal of Political Economy*, 100, 992–1026. [650]

Bolton, Patrick and Christopher Harris (1999), "Strategic experimentation." *Econometrica*, 67, 349–374. [650]

Brémaud, Pierre (1999), *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues.* Springer-Verlag, New York. [674]

Bulow, Jeremy and Paul Klemperer (1994), "Rational frenzies and crashes." *Journal of Political Economy*, 102, 1–23. [652]

Callander, Steven and Johannes Hörner (2009), "The wisdom of the minority." *Journal of Economic Theory*, 144, 1421–1439. [650]

Çelen, Boğaçhan and Shachar Kariv (2004), "Observational learning under imperfect information." *Games and Economic Behavior*, 47, 72–86. [650]

Chamley, Christophe P. (2004), *Rational Herds: Economic Models of Social Learning*. Cambridge University Press. [650]

Chaudhry, Mohan L. and Umesh C. Gupta (1996), "Performance analysis of the discretetime GI/Geom/1/N queue." *Journal of Applied Probability*, 33, 239–255. [650, 652]

Debo, Laurens G., Christine Parlour, and Uday Rajan (2012), "Signaling quality via queues." *Management Science*, 58, 876–891. [650, 651]

Debo, Laurens G., Uday Rajan, and Senthil Veeraraghavan (2018), "Signaling quality via long lines and uninformative prices." Unpublished paper, SSRN 2033909. [650, 651]

Eyster, Erik, Andrea Galeotti, Navin Kartik, and Matthew Rabin (2014), "Congested observational learning." *Games and Economic Behavior*, 87, 519–538. [650, 651]

Hassin, Refael (1985), "On the optimality of first come last served queues." *Econometrica*, 53, 201–202. [650]

Hassin, Refael and Moshe Haviv (1995), "Equilibrium strategies for queues with impatient customers." *Operations Research Letters*, 17, 41–45. [651]

Hassin, Refael and Moshe Haviv (2003), *To Queue or not to Queue: Equilibrium Behavior in Queuing Systems*. Kluwer Academic Publishers, Dordrecht, The Netherlands. [651]

Haviv, Moshe and Ya'acov Ritov (2001), "Homogeneous customers renege from invisible queues at random times under deteriorating waiting conditions." *Queueing Systems*, 38, 495–508. [651]

Herrera, Helios and Johannes Hörner (2013), "Biased social learning." *Games and Economic Behavior*, 80, 131–146. [651]

Keller, Godfrey, Sven Rady, and Martin Cripps (2005), "Strategic experimentation with exponential bandits." *Econometrica*, 73, 39–68. [650]

Mandelbaum, Avishai and Nahum Shimkin (2000), "A model for rational abandonments from invisible queues." *Queueing Systems*, 36, 141–173. [651]

Murto, Pauli and Juuso Välimäki (2011), "Learning and information aggregation in an exit game." *The Review of Economic Studies*, 78, 1426–1461. [650, 652]

Naor, Pinhas (1969), "The regulation of queue size by levying tolls." *Econometrica*, 37, 15–24. [650, 653, 654]

Percus, Ora E. and Jerome K. Percus (1990), "Elementary properties of clock-regulated queues." *SIAM Journal Applied Mathematics*, 50, 1166–1175. [650]

Shimkin, Nahum and Avishai Mandelbaum (2004), "Rational abandonment from telequeues: Nonlinear waiting costs with heterogeneous preferences." *Queueing Systems*, 47, 117–146. [651]

Smith, Lones and Peter Sørensen (2000), "Pathological outcomes of observational learning." *Econometrica*, 68, 371–398. [650]

Smith, Lones and Peter Norman Sorensen (2013), "Rational social learning by random sampling." Unpublished paper, SSRN 1138095. [650]

Strulovici, Bruno (2010), "Learning while voting: Determinants of collective experimentation." *Econometrica*, 78, 933–971. [650]

Thomas, Caroline D. (2018), "Strategic experimentation with congestion." Unpublished paper, SRRN 3266914. [650]

Toxvaerd, Flavio (2008), "Strategic merger waves: A theory of musical chairs." *Journal of Economic Theory*, 140, 1–26. [652]

Wolff, Ronald W. (1982), "Poisson arrivals see time averages." *Operations Research*, 30, 223–231. [652]

Co-editor Johannes Hörner handled this manuscript.

Manuscript received 3 March, 2014; final version accepted 29 July, 2018; available online 2 August, 2018.