

Strategic Experimentation in Queues

Martin W. Cripps* and Caroline D. Thomas†

February 25, 2014

Abstract

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 which 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 only observe the evolution of the queue after they arrive, thus each agent solves the experimentation problem of 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. There is also a negative payoff externality as those at the front of the queue delay the service of those at the back.

Journal of Economic Literature Classification Number: C72, C73.

Keywords: Experimentation, Bandit Problems, Social Learning, Herding, Queues.

*University College London. I am grateful to the Cowles Foundation for its hospitality while part of this work was undertaken.

†University of Texas at Austin. Support from Deutsche Bank through IAS Princeton is gratefully acknowledged.

1 Introduction

We study a game of strategic experimentation that has both payoff and information externalities. A sequence of individuals arrive over time and join a queue for service. This queue grows at each new arrival and shrinks if service occurs or if an individual decides to stop waiting and leave. Individuals arrive uncertain about whether service occurs, because in the bad state of the world there is no service, but once in the line can observe the service events as well as the behaviour of all other agents in the queue.

As she waits in line without observing service, an individual revises downwards the likelihood she attributes to service ever occurring. This is the usual private learning that occurs in strategic experimentation models. In the standard exponential bandit problem a single agent would decide how long to wait for a reward/service before giving up and taking an outside option. This aspect is also present in our model. Additionally, the behaviour of other agents in this game is itself a source of information. This social learning takes two different forms. Once in the queue, individuals learn from the behaviour (leave or keep queueing) of those ahead of her in the queue. For instance, observing an agent ahead of her leaving the queue is bad news about the state of the world. Social learning also occurs when the individual arrives at the queue. The service state determines the stochastic process followed by the queue lengths, so the length of the queue when she arrives is informative about this state.

Our main results are as follows. We find a class of strategies that combine herding and experimentation in a natural way. We establish the existence of a symmetric equilibrium in such strategies when agents are sufficiently patient, and describe the equilibria obtaining when agents are less patient. Depending on the discount factor, equilibria can take two qualitatively different forms. When agents are sufficiently patient they are willing to let queues grow very long. These queues can be very informative, and certain queue lengths *perfectly reveal* the state of the world to new arrivals. This is not the case when agents are less patient. Agents are then unwilling to let the queues grow long and no queue length can perfectly reveal the state.

The equilibrium properties also depend on the rate at which service occurs in the good state. If the (good state) rate of service is greater than the arrival rate (so that queues tend to empty out), we find that in equilibrium there are positive spillovers in the experimentation decisions. That is, the individuals in the queue tend to experiment for longer than they would have in the corresponding single-agent decision problem. This is because short queues signal a good state of the world and cause an individual arriving first in line to revise upwards her belief in the good state. This increased waiting time is good for social learning, because many other agents will be able to benefit from the first in line's experience. However, the tendency of queues to empty out is bad for social memory. Every time the queue clears the social memory is reset and individuals have to re-learn what past generations may already have learnt. Ultimately the equilibrium in this case does not get close to efficiency.

In contrast if (in the good state) the arrival rate is greater than the service rate, there are negative spillovers in the experimentation decisions. Individuals early in line experiment

less than they would have in the single-agent problem, which is bad for social learning. However, the tendency of queues to fill up in this case implies that once the state is known to be good the queue will tend to persist for long periods of time without clearing. Social memory is therefore excellent. As a result the equilibrium behavior can be close to efficient.

1.1 Related Literature

The model we study formalizes a problem that arises in many contexts. 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 queues a good sign because they indicate a high service rate, or a bad sign because informed individuals know not to queue but go straight to the outside option? This problem also arises in many non-economic situations¹ (queueing for service in computer and communication networks, pipeline scheduling).

There is a vast literature on strategic behaviour in queues (see Hassin and Haviv (2003) for a summary) and it is well known that in queues operating under a first-in-first-out (FIFO) regime an individual who decides to join the queue imposes a negative payoff externality on those behind her (see Hassin (1985)). Research most closely related to ours considers the question of herding and social learning (Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992), Smith and Sørensen (2000)). In the context of queues Debo, Parlour, and Rajan (2012) consider a model in which the length of a queue reveals agents' private information about the quality of a product, and explore a firm's incentive to manipulate the service rate or (in Debo, Rajan, and Veeraraghavan (2012)) prices. Eyster, Galeotti, Kartik, and Rabin (2013) 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. In all these models, queues serve to add a cost to herding and all learning is done prior to the individuals' decision whether to join the queue or not. Once an individual has made this decision, she cannot revoke it, and there is no further learning, public or private.

Strategic experimentation with information externalities has been widely studied (Bolton and Harris (1999), Keller, Rady, and Cripps (2005), Murto and Välimäki (2011)) and there is a recent literature on experimentation with direct payoff externalities (Strulovici (2010), Thomas (2013)). 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 studied.

In the queue setting we are also able to make a distinction between social *learning* and social *memory*. The current size of the queue and the current behaviour of those in the queue generate an informational externality giving rise to social learning. However, in our model there are events that can destroy the accumulated knowledge of everyone in the queue. This can happen if the entire queue is served and clears, or if all agents in the queue leave en masse. After such events, the new arrivals find themselves in a world where there are no agents to learn from and there is no record of what went on before. The social record was obliterated. These events can happen with positive probability so that

¹See Percus and Percus (1990) or Chaudhry and Gupta (1996) for examples.

in this game learning never stops. We define social memory to be the *persistence* of the social learning, and the frequency of the events thus resetting social learning determines the social memory. Similar issues, although in a different context, have been discussed in Herrera and Hörner (2013).

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).

The layout of the paper is as follows. In Section 2 we set up our discrete-time queuing model. In Section 3 we describe each individual’s experimentation problem. Section 4 discusses the inefficiencies that arise in this game. Section 5 provides the main result of the paper. We establish the existence of a symmetric equilibrium when individuals are sufficiently patient. We describe the social learning process as a function of the arrival rate of rewards in the good state and of agents’ impatience. We discuss social memory in Section 6 and directions for further research in Section 7.

2 The Model

We start by describing the queue. A single-server queuing model is described by two processes: one governs the arrival of individuals, the other determines when they are served. Below we will describe the details of the arrival process and the two different service processes that may occur in our model, depending on the server state. We then describe the agents’ beliefs about the server state, and finally their payoffs.

Time is discrete and indexed by $\tau = 0, 1, \dots$. At each date τ , we distinguish three separate stages: *Service*, *Exit*, *Arrival*. The *S, E, A* stages collapse separate events in the queueing process so that they occur at the same calendar time. At any date τ the *S, E, A* stages proceed as follows:

Service: This is first stage in each period. Let k_τ denote the number of individuals that are served at date τ . There are two possible states of service: {good, bad}. In the bad state no individual is ever served: $k_\tau = 0$ for $\tau = 0, 1, 2, \dots$. In the good state k_τ , is an *i.i.d.* random variable with a geometric distribution²: $\Pr(k_\tau = x) = (1 - \alpha)\alpha^x$ for $x = 0, 1, \dots$ and $\alpha \in (0, 1)$.

Exit: In the second stage individuals have the opportunity to leave the queue. (This is also called “reneging” on the decision to queue.) Any exit is observed by all individuals who are currently in the queue. This stage is only concluded when no individual remaining in the queue wishes to exit, so there is the opportunity for multiple rounds of exit at this stage if this is desired by the agents. The opportunity to exit occurs immediately after service occurs. She would not find it optimal to leave the queue at any other stage of date τ , because this is the point at which she has just learnt something about the service rate.

Arrival: At the final stage one new individual arrives. The new individual at date τ can choose to join the queue or to “balk”, that is, to exit immediately and not join the

²This discrete-time geometric distribution service model for queues is widely used to model computer communication system: see for example Chaudhry and Gupta (1996).

queue. Once this stage has occurred the game moves to the next time period and the S, E, A sequence is repeated.

Notice that in this model individuals must wait in line for at least one period before they have the opportunity to be served. Also, that the average rate at which individuals are served in the good state is $\alpha/(1 - \alpha)$. If $\alpha > 1/2$, the average service rate is greater than the arrival rate and queues tend to empty whereas if $\alpha < 1/2$ queues tend to grow.

The individual who arrives at date τ is uncertain about the state of the server and about the current date. She holds a prior belief on both. Let $\mu \in (0, 1)$ denote each individual's prior belief that the server is in the good state and each individual assigns prior probability $\nu(1 - \nu)^\tau$ to having arrived into the system at date τ .³ Once she observes the number of individuals already queueing for service she will revise these beliefs and can then (as specified above) choose to join the queue or to balk.

We assume that any individual who is served receives the payoff of w . Any individual who exits, either initially or after waiting for some time (balks or reneges), receives a payoff of 1. While individuals wait in line they receive a flow payoff of zero and discount one unit of calendar time by the factor $\delta < 1$. Finally, we will assume that $\delta w > 1$ so it is optimal to wait for service if the server is known to be good.

3 Optimal Experimentation in Queues

We now describe the solutions to two single-agent optimization problems that are components of our game. The first question is: when should an individual join a queue if she knows the server state is good? This question arises because, even in the good state, it may take significant time for a long queue to be served. So, a new arrival may prefer balking and immediately taking the outside option to waiting in line a long time. The answer to this question will determine the maximum individually rational queue length \mathcal{M} . The second problem is one of optimal experimentation: How long should an individual who is n^{th} in line wait without observing service before reneging on the queue and taking the outside option? In this section we treat this question as a private-learning problem and assume that the individual does not learn from the actions of others in the queue.

3.1 The Maximal Individually Rational Queue Length: \mathcal{M}

First we evaluate the value of being n^{th} in line when the server is known to be good. We will introduce the parameter ψ , which represents the congestion cost of being behind another individual in a queue. Once we have determined the value of being n^{th} in line, we can compare this with the value of taking the outside option and determine M , the maximal individually rational queue length at a server known to be good.

Let V_n denote the expected payoff of an individual who at date τ 's service opportunity

³Our results will apply to the case where ν is sufficiently small and this prior is sufficiently diffuse. The parametrization of this prior is not significant it will be simply a way of ensuring that this prior is diffuse enough.

has just observed service and is now n^{th} in the queue. This satisfies the recursion

$$V_n = (1 - \alpha)\delta V_n + (1 - \alpha)\alpha\delta V_{n-1} + \cdots + (1 - \alpha)\alpha^{n-1}\delta V_1 + \alpha^n\delta w.$$

(At the next service opportunity exactly $x = 0, 1, \dots, n - 1$ individuals are served with probability $(1 - \alpha)\alpha^x$ and the n^{th} in line moves up to the $n - x^{\text{th}}$ position. With probability α^n at least n individuals are served, including the n^{th} in line.) As $V_1 = (1 - \alpha)\delta V_1 + \delta\alpha w$, we can solve iteratively to find

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

This expression serves several purposes. First, it will be an input into the calculation for the optimal time to wait when the server state is unknown. Second, each additional person queuing in front of her discounts an individual's payoff by the factor $\psi < 1$. The parameter ψ can be understood to capture the *congestion cost* imposed by any individual on those behind her in the queue. This congestion cost is mitigated as the service rate, α , increases and it entirely disappears as α , and so ψ , approach one. In contrast when service is slow the congestion costs become extreme.

We now describe queues where congestion cost are large enough to make balking preferable to waiting. When the server is known to be good, if $\psi^n\delta w > 1$ an individual prefers joining the queue at the n^{th} position to balking. We define $\mathcal{M} \geq 0$ to be the largest integer such that

$$(2) \quad \psi^{\mathcal{M}+1}\delta w < 1 \leq \psi^{\mathcal{M}}\delta w.$$

\mathcal{M} is the longest the queue can ever get. It depends⁴ only on the parameters α , δ and w of our model. Our assumption on w ensures \mathcal{M} is positive. Notice that infinitely long lines are possible as congestion costs vanish, that is, as the individuals become more patient ($\delta \rightarrow 1$) or as the service rate increases ($\alpha \rightarrow 1$).

3.2 The n^{th} in Line's Experimentation

We now turn to an individual's private learning, or *experimentation*. This learning is based only on her observations of the server (in)activity as she waits in line. Hence, in this section we will assume that she learns nothing from the actions of other agents – there are no informational externalities. We maintain the assumption that a player must wait for those in front of her to either be served or to renege on the queue, before she can be served. This generalizes the usual bandit problem to one where arrival of good news does not immediately generate a reward; the reward (service) arrives at some random time in the future.

We want to determine the length of time for which an uninformed agent, who is n^{th} in line, will optimally wait to learn the queue state. Our first step is to evaluate the expected payoff of the individual who is n^{th} in line conditional on at least one individual

⁴For notational convenience we will not make the dependence on these parameters explicit.

being served. This is essentially the n^{th} in line's expected payoff conditional on service occurring at the current service opportunity. A simple calculation (and a substitution from the recursion above for V_{n-1}) gives this value:

$$(3) \quad (1 - \alpha) (V_{n-1} + \alpha V_{n-2} + \cdots + \alpha^{n-2} V_1) + \alpha^{n-1} w = V_{n-1} / \delta = \psi^{n-1} w$$

Thus the n^{th} in line expects to get a payoff $\psi^{n-1} w$ if the server is revealed to be good at the current service opportunity. For $n = 1$ the expected payoff from service having occurred is w , as one would expect, and in general the expected payoff conditional on service having occurred is proportional to the value (V_{n-1}) of being $n - 1^{\text{st}}$ in a good queue.

Given this preliminary calculation, we can now describe the payoff, $U_n(m, \mu_n^0)$, of an individual who arrives as the n^{th} in line, has belief $\mu_n^0 > 0$ that the server is in the good state⁵ 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 m periods of server inactivity. The details of $U_n(m, \mu_n^0)$ can be explained as follows. First, the individual expects to observe no service over m periods with probability $1 - \mu_n^0 + \mu_n^0(1 - \alpha)^m$. If service occurs in any of the m periods, her expected payoff is given by (3).

$$(4) \quad U_n(m, \mu_n^0) := (1 - \mu_n^0 + \mu_n^0(1 - \alpha)^m) \delta^m + \psi^{n-1} w \mu_n^0 \sum_{s=1}^m \delta^s \alpha (1 - \alpha)^{s-1},$$

$$(5) \quad = (1 - \mu_n^0) \delta^m + \mu_n^0 \psi^n w \delta - \mu_n^0 \delta^m (1 - \alpha)^m (\psi^n \delta w - 1).$$

The three terms on the right of (5) represent: her payoff from taking her outside option when the state is bad, her payoff from always being served 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 individual who is n^{th} in line will solve the problem $\max_{m \geq 0} U_n(m, \mu_n^0)$. Her optimal behavior could be described in terms of a cutoff posterior $\underline{\mu}_n$ at which she should renege on the queue, or in terms of the number $\mathcal{N}(n, \mu_n^0) := \arg \max_{m \geq 0} U_n(m, \mu_n^0)$ of unsuccessful service events she should observe before renegeing. The result below describes both.

Proposition 1 *There exists a solution, m^* , to the problem $\max_{m \geq 0} U_n(m, \mu_n^0)$. The optimal action, m^* is unique for a.e. $\mu_n^0 \in (0, 1)$. At beliefs where the solution is not unique, m^* and $m^* + 1$ are both optimal. The value m^* satisfies*

$$(6) \quad m^* = \mathcal{N}(\mu_n^0, n) := \left\lceil (\log(1 - \alpha))^{-1} \log \left(\frac{1 - \mu_n^0}{\mu_n^0} \frac{\psi(1 - \delta)}{\alpha(\psi^n \delta w - 1)} \right) \right\rceil_+,$$

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

$$(7) \quad \underline{\mu}_n := \frac{1 - \delta}{\delta \alpha (\psi^{n-1} w - 1)}.$$

The proof is given in Appendix A.1.

⁵In Section 5.3 we describe how this belief is obtained.

4 The Cooperative Problem

We now describe how a team of individuals can act to maximize average social welfare.⁶ This is the benchmark for the later results when the individuals act strategically. We will show that, when acting as a team, the individuals can achieve the ex-post first-best average social welfare. That is, there exists a strategy that allows the team to achieve the first-best average payoff when the queue is good *and* the first-best average payoff when the queue is bad! This strategy will generate a shorter maximum queue than that observed in the equilibrium of section 5.

If the members of the team can communicate information to later arrivals, the team problem is trivial, because in the good state service will eventually be observed and all later arrivals will be told the state. We will, therefore, consider efficiency where the members of the team cannot share information with each other: each member of the team is constrained to follow a strategy that depends only on their information. The objective of the team is to maximise the long-run average utility of all team members. There is no discounting of team-members' utilities and we consider a Utilitarian social welfare criterion.

First let's consider the cooperative optimum when the state of the server is known. In the bad state the cooperative solution would be for every arrival to balk immediately. This corresponds to the individually optimal behaviour. In the good state, however, individually and socially optimal behaviour differ: In the good state, it is individually optimal for individuals to join queues shorter than \mathcal{M} (where \mathcal{M} is defined by condition (2) above) and to balk at queues of lengths \mathcal{M} or greater. This is not the cooperative optimum, however. The individual who is \mathcal{M}^{th} in line is close to indifferent between waiting and immediate exit; the private benefits and costs of standing in line are close to being equal. However, waiting in line imposes additional costs on the team, because it delays the expected service time of all later arrivals (until the line empties). Hence, the privately optimal decision of the \mathcal{M}^{th} in line has additional social costs that are not incorporated into her private decision. It follows that the socially optimal line length, M^\dagger , in the good state satisfies $M^\dagger \leq \mathcal{M}$ and in general this inequality is strict. Individually optimal queues at good servers are in general inefficiently long.⁷

Now let us return to the case where the server state is unknown. The strategy we propose for each individual is: join the queue and wait until served if there are less than M^\dagger individuals in line; balk and take the outside option immediately if there are M^\dagger or more in the line. This strategy achieves the first-best cooperative payoff in the good state of the world. In the bad state of the world, the first M^\dagger arrivals will join the queue and, thereafter, no-one will. Averaged over infinitely many individuals, this imposes a vanishingly small social cost and the average payoff in the bad state of the world is 1. The policy is, therefore, also ex-post optimal given the team's Utilitarian objectives.⁸

⁶ In a queuing model this problem was studied by Hassin (1985), in the social learning setting Smith and Sørensen (2009) and in the bandit setting by Bolton and Harris (1999); for example.

⁷The exact value of M^\dagger is calculated in Appendix B.

⁸ In a more general model where service was provided in the bad state with a reduced probability the ex-post optimum would not be obtainable. This is because following the optimal good-state strategy in the bad state would impose increased waiting on an infinite set of agents. We conjecture that the optimal strategy for the team would be to exit at a threshold that was strictly between the exit threshold in the

5 Equilibrium of the Queuing Game

We now construct a symmetric equilibrium of this game. The first step in this process is to propose a strategy that will be played by each agent. Then, for each server state we determine the stationary distribution of queue lengths induced by this strategy. These stationary distributions determine an agent's posterior upon arriving at the queue at a given position and observing the current queue length. Finally, we verify that the strategy assumed at the first step is indeed optimal given this posterior and the other learning that occurs in this game. Each part of this section deals with one of the four steps just described.

5.1 Strategies

An individual's strategy has two parts: It prescribes her behaviour when she arrives at the queue (whether to join the queue or to balk), and her behaviour once she has joined the queue (when to stay or when to renege).

Definition 1 *The strategy $\sigma^*(q, N, M)$:*

- *An individual joins any queue as long as she is at most M^{th} in line.*
- *If there are agents ahead of her in the queue when she joins it, then she reneges on the queue if and only if the first in line reneges.*
- *If there are no agents ahead of her in the queue when she joins it, then she reneges after N unsuccessful service opportunities with probability $q \in [0, 1]$. With probability $1 - q$ she experiments for one more period and reneges if not served at the $N + 1^{\text{st}}$ service opportunity.*

The probability $q \in [0, 1]$ and the non-negative integers N and M are parameters of σ^* .

According to this strategy,⁹ no individual other than the first in line autonomously reneges on the queue. For $N < M$ this means that agents will continue to join the queue as long as the first in line experiments. If an individual joins a queue no longer than N , she does not know whether the individual 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, the agent joining a queue may not know whether the first in line has already observed service or is still experimenting.

If $N < M$ the length of the queue can also reveal the first in line's information. If the first in line reneges after N periods of experimentation, all those behind her infer that she has not yet observed service. If the first in line does not renege after N periods of experimentation, all those behind her learn with certainty that the first in line has previously observed service. They are now certain that the server is in the good state and

good state and the bad state.

⁹In a slight abuse of notation, we will use σ^* interchangeably to denote an individual's strategy or the symmetric strategy profile in which every individual uses the strategy σ^* .

will stay in the queue until served. Therefore, an individual arriving at a queue in n^{th} position, where $N < n \leq M + 1$, can be certain that the server is in the good state. We say that the strategy profile σ^* exhibits *perfect revelation* when $N < M$.¹⁰

If on the other hand $M \leq N$, the queue will not exceed length M even as the first in line continues to experiment. All agents queuing behind the first in line learn that the server is in the good state if the first in line doesn't renege after N unsuccessful service events, but even in that case the queue never grows longer than M . So while the position n at which an agent arrives at the queue remains informative about the server state, there exists no n that perfectly reveals the server state. We say that the strategy profile σ^* exhibits *imperfect revelation* when $M \leq N$.¹¹

We will establish that a strategy profile σ^* with perfect revelation ($N < M$) constitutes a perfect Bayesian equilibrium of the game, provided agents are sufficiently patient and have a sufficiently diffuse prior on the calendar date at which they enter the system. This is summarised in the proposition below. The result holds for all $\alpha \in (0, 1)$. When δ is sufficiently large, it is optimal for the first in line to engage in some experimentation. Moreover the cost of congestion is sufficiently small for agents to be willing to join long queues. Both of these elements ensure that there exists a PBE with perfect revelation.

Proposition 2 *Given $\mu \in (0, 1)$ and $\alpha \in (0, 1)$ there exists $\underline{\delta} < 1$ such that for all $\delta > \underline{\delta}$ and all $\nu < \bar{\nu}(\delta)$ there exists a strategy $\sigma^*(q^*, N^*, M^*)$ with perfect revelation that constitutes a symmetric perfect Bayesian equilibrium of this game.*

This proposition is proved by Lemmas 5 and 7 below. Equilibria with imperfect revelation exist for lower values of δ . While we don't prove this result formally, we give an intuition for it in Section 5.5 where we give examples of such equilibria.

5.2 Stationary Distributions of Queue Lengths

The inference an individual draws from the queue length she observes upon arrival is described in Section 5.3. It depends on the distributions of possible queue lengths conditional on the server state and given that all individuals use the strategy σ^* described above. We describe the stationary distributions of queue lengths in this section.

We will consider the stochastic process followed by the queue length *at the start of the arrival stage* of each date $\tau \in \mathbb{Z}_+$. This choice of timing is critical in understanding the distributions below. We say that the queue has length n at date τ if the individual arriving in the system at date τ arrives in the queue at the n^{th} position—even if that individual then balks. There are two discrete-time Markov processes to consider: one that arises if the server is in the good state and the other if the server is in the bad state. We now derive the stationary distributions of these processes.

Conditional on the server being in the bad state, the queue length (at the start of the arrival stage) follows an almost deterministic process. Let w_n ($n = 1, 2, \dots, M + 1$) denote the stationary probability of arriving at the queue at the n^{th} position, conditional on the

¹⁰For $q < 1$ perfect revelation requires $N + 1 < M$.

¹¹For $q < 1$ the strategy profile with $M = N + 1$ also exhibits imperfect revelation.

server being in the bad state. If $N < M$ the queue grows by one individual each period and then shrinks to length 1 in the period after reaching length N with probability q , or with probability $1 - q$ grows to length $N + 1$ and then shrinks to length 1. Thus, in the bad state there is ergodic probability $1/(N + 1 - q)$ of arriving n^{th} in line for $n = 1, 2, \dots, N$, probability $(1 - q)/(N + 1 - q)$ for $n = N + 1$, and zero probability of arriving at the queue at any other position. If $M \leq N$ the queue will grow to size M and then stay at that size for a further $N - M$ periods with probability q or a further $N + 1 - M$ periods with probability $1 - q$ before shrinking to unity. Thus there is ergodic probability $1/(N + 1 - q)$ of arriving at a queue n^{th} in line for $n \leq M$ and ergodic probability $(N - M + 1 - q)/(N + 1 - q)$ of arriving at the $M + 1^{\text{st}}$ position (and balking). These values are summarized in Proposition 3 below.

Conditional on the server being in the good state, the process governing the evolution of the queue (at the start of the arrival stage) is more complex. If the agents use the strategy $\sigma^*(q, N, M)$, the queue length at the end of each period follows a stochastic process: sometimes service occurs and shrinks the queue and other times it does not; sometimes the first in line reneges and all the others in line follow her. This stochastic process is a Markov chain, provided the state of the process is defined to be the position in the queue at which the latest individual arrives *and* whether or not the first in line knows that the server is in the good state. There are at most $M + N + 2$ states for this process: arrival at positions $1, 2, 3, \dots, M + 1$ and the first in line knows the server is in the good state; arrival at positions $1, 2, \dots, N + 1$ and the first in line is uncertain. The process governing the queue length, defined by the strategy above and the service process, has finite states and is irreducible so it must admit a unique stationary measure.

We define y_n ($n = 1, 2, \dots, M + 1$) to be the stationary probability of arriving n^{th} in line at the arrival stage of date $\tau \in \mathbb{Z}_+$, conditional on the server being in the good state. We also define z_n ($n = 1, 2, \dots, N + 1$) to be the stationary probability that the first in line has not observed a service event and that the individual arriving at date $\tau \in \mathbb{Z}_+$ arrives at the n^{th} position, conditional on the server being in the good state.¹² These values are characterised in the following proposition.

Proposition 3 *Assume that $\alpha \neq 1/2$ and $\alpha \neq \alpha_N^*$ where α_N^* solves $(1 - \alpha)^{N+1} = 1 - 2\alpha$. If the agents follow the strategy $\sigma^*(q, N, M)$, then conditional on the server being in the good state the unique stationary distribution satisfies $z_n = (1 - \alpha)^{n-1}y_1$ for $n = 1, 2, \dots, N + 1$. If $N < M$, then*

$$(8) \quad y_n = B \begin{cases} \phi^{n-1} - k_N, & n = 1, 2, \dots, N; \\ \phi^{n-1} - k_N \frac{q + (1-q)\alpha\phi^2}{\alpha\phi^2 + q(1-\alpha)}, & n = N + 1; \\ \phi^{n-1} - k_N \frac{1+q\phi}{\phi^2(\phi+q)} \phi^{n-N}, & n = N + 2, \dots, M + 1; \end{cases}$$

$$(9) \quad B^{-1} = \frac{1 - \phi^N}{1 - \phi} - Nk_N + \frac{\phi^N - \phi^{M+1}}{1 - \phi} \left(1 - \frac{k_N}{\phi^{N+1}} \frac{1 + q\phi}{\phi + q} \right) + (1 - q)k_N \frac{1 - \phi^2}{\phi(\phi + q)}.$$

¹²The remaining part of the stationary distribution can be found by taking the difference $y_n - z_n$. Notice that an individual cannot have just arrived at the first position in the queue and know that the server is good. This is why the state where an individual arrives at the queue in first position and knows that the server is in the good state has zero measure, and $y_1 = z_1$.

If $M \leq N$, then

$$(10) \quad y_n = B \begin{cases} \phi^{n-1} - k_N, & n = 1, 2, \dots, M; \\ \phi^M - \frac{k_N}{\phi}, & n = M + 1; \end{cases}$$

$$(11) \quad B^{-1} = \frac{1 - \phi^{M+1}}{1 - \phi} - k_N(M + \phi^{-1}).$$

In both cases: $\phi := (1 - \alpha)/\alpha$; $k_N := \alpha(\phi + q)(1 - \alpha)^{N+1}/[\alpha(\phi + q)(1 - \alpha)^{N+1} + 2\alpha - 1]$. Finally, conditional on the server being in the bad state, the unique stationary distribution satisfies, for $N < M$:


$$(12) \quad w_n = \begin{cases} \frac{1}{N+1-q}, & n = 1, 2, \dots, N; \\ \frac{1-q}{N+1-q}, & n = N + 1; \\ 0, & n = N + 2, \dots, M + 1; \end{cases}$$

and for $M \leq N$:

$$(13) \quad w_n = \begin{cases} \frac{1}{N+1-q}, & n = 1, 2, \dots, M; \\ \frac{N-M+1-q}{N+1-q}, & n = M + 1. \end{cases}$$

The proof of this result is given in Appendix A.2. In the statement of Proposition 3 there are two excluded values of α . In both cases the stationary distribution exists, but it has a different functional form from the ones given above. We will describe these distributions later in this section. For most values of α , this stationary distribution admits three qualitatively different forms:¹³

Decreasing when $\alpha > 1/2$: In this case service is faster than arrivals, so shorter queues are more likely than longer ones. 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 faster than exponential decline when $M \leq N$, and for the values $n = 1, \dots, N$ when $N < M$. The jump down between $n = N$ and $n = N + 1$ occurs because such a transition is only possible if the first in line knows that the server is in the good state. Similarly for the jump between $N + 1$ and $N + 2$ when $q < 1$. For $n = N + 2, \dots, M$, the distribution declines exponentially.



¹³All 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.

Figure 1: The stationary measure of the queue length conditional on the server being in the good state with $\alpha > 1/2$ under perfect revelation (left panel) and imperfect revelation (right panel).

U-Shaped when $\alpha < 1/2$ and $k_N > 1$: This occurs for $N < M$ when α takes values in the interval $(\alpha_N^*, 1/2)$. That interval vanishes ($\alpha_N^* \rightarrow 1/2$) as $N \rightarrow \infty$. For these values of α , service is slower than arrivals so that, unconditionally, longer queues are more likely than shorter ones. However the effect of slow service is dominated by the renewal effect when $M \leq N$ and for the values $n = 1, \dots, N$ when $N < M$. Therefore, conditional on the first in line being uninformed, shorter queues are more likely than longer ones and the stationary distribution is declining with n . In contrast, once the queue grows longer than N it tends to fill up to length M and stay there for some time. So the stationary distribution jumps down at $N + 1$ and $N + 2$ and then increases over the range $N + 2 \leq n \leq M$, as illustrated in Figure 3 below.




Figure 2: The stationary measure of the queue length conditional on the server being in the good state with $\alpha_N^* < \alpha < 1/2$ under perfect revelation (left panel) and imperfect revelation (right panel).

Increasing when $k_N < 0$: This occurs for $\alpha \in (0, \alpha_N^*)$. In this case service is so slow that it dominates the renewal effect. The stationary measure is therefore increasing over its entire support. Notice that as M increases without bounds, y_1 tends to zero. Intuitively: If the queue at a good server is most likely to be infinitely long, arriving at the first position in line makes a individual almost certain that the server is in the bad state.




Figure 3: The stationary measure of the queue length conditional on the server being in the good state with $\alpha < \alpha_N^*$ under perfect revelation (left panel) and imperfect revelation (right panel).

Excluded values $\alpha = 1/2$ and $\alpha = \alpha_N^*$: For $\alpha = 1/2$, the exact analytical form of the stationary distribution is derived in Appendix A.2. For $N < M$ it is linearly decreasing in n for $1 \leq n \leq N$, has a downward step at $n = N + 1$ and $n = N + 2$, and is constant for $n \geq N + 2$. For $M \leq N$ it is linearly decreasing in n for $1 \leq n \leq M$ and has a downward step at $n = M + 1$.

For $\alpha = \alpha_N^*$ and $N < M$ the stationary measure is uniform for $n \leq N$, has a downward step at $n = N + 1$ and $n = N + 2$, and is increasing for $n \geq N + 2$. For $M \leq N$ it is uniform for $1 \leq n \leq M$ and has a downward step at $n = M + 1$. These cases are illustrated below.




Figure 4: The stationary measure of the queue length conditional on the server being in the good state when α takes the values α_N^* and $1/2$ under perfect revelation (first two panels) and imperfect revelation (last two panels).

We now establish a bound on the rate at which the Markov process followed by the queues in the good state converges to the stationary distributions defined in Proposition 3. This result will be of use in the next section when we argue that posteriors based on these stationary distributions are a good approximation to the agents' true posteriors when they arrive at a queue of a given length. There are at most $M + N + 2$ states for this Markov process that is: arrival at positions $1, 2, 3, \dots, M + 1$ and the first in line knows that the server is in the good state; arrival at positions $1, 2, \dots, N + 1$ and the first in line is uncertain about the server state. Let $S := \{1, 2, \dots, M + 1\} \cup \{1, \dots, N\}$ denote this state space and $\zeta \in \Delta(S)$ a generic probability measure on S . The initial condition (at the arrival state of date $\tau = 0$) is that one (uniformed) individual arrives at the first position in line. We will denote its measure as $\zeta^0 \in \Delta(S)$. The strategy $\sigma^*(q, N, M)$ together with the service process determine a probability distribution for queue lengths at all future dates $\tau = 1, 2, 3, \dots$, which we will denote $\zeta^\tau \in \Delta(S)$. Finally we let $\bar{\zeta} \in \Delta(S)$ denote the stationary measure described in Proposition 3. In the lemma below we give a rate of convergence result for this process. That is, we bound the distance between ζ^τ and $\bar{\zeta}$, where $\|\cdot\|$ denotes the total variation norm.

Lemma 1 *If the server is in the good state, then $\|\zeta^\tau - \bar{\zeta}\| < (1 - \alpha^M)^\tau$ for all $\tau > 0$.*

The proof of this lemma can be found in the Appendix A.4. It is intuitive once one appreciates that at each date there is a probability of at least α^M that all the individuals in the queue are served (no matter how long the queue is). Once all individuals are served, the queue reverts to the state in which an individual arrives at the queue at the first position and is hence uninformed about the server state. This renewal (or coupling) rate bounds the rate of convergence to the stationary measure.

5.3 Posteriors and Inference on Queue Lengths

In this section we describe the properties of the agents' posterior beliefs when the strategy $\sigma^*(q, N, M)$ is used. There are two sources of information and, therefore, two kinds of learning in this model. We refer to the inference an agent draws from her own observation of server activity as private learning, or experimentation, and distinguish it from social learning, where an agent draws inference from the actions of other agents. At our equilibrium social learning occurs when agents arrive in the system and observe the queue and when they observe the first in line's behaviour after N or $N + 1$ unsuccessful service events.

We begin this section by giving two expressions, (14) and (15), for the updated beliefs the individual *would have* if she were certain that the process governing the queue length is in the stationary regime. (The social learning that an individual performs upon arriving at the queue at the n^{th} position, without incorporating any prior information the individual may have about the calendar date.) Our first result in this section, Lemma 2, shows that these expressions are arbitrarily good approximations to the individual's true posterior (incorporating her belief about the date) provided the parameter ν is sufficiently close to zero.

The second result describes the relationship between private and social learning. In Lemma 3 we show that an agent later in the queue is *always* more optimistic than those ahead of her (conditional on no agent having actually observed service). In our final result of this section, Lemma 4, we will describe how agents' social learning varies with the parameters N of the strategy $\sigma^*(q, N, M)$ and α of the arrival process.

Consider an individual arriving at the n^{th} position in the queue. Her observation of the queue length provides information about the state of the server. We use μ_n^0 to denote her updated prior, that is, the individual's posterior belief that the server is in the good state conditional on arriving at the queue at the n^{th} position. We use $\bar{\mu}_n^0$ to denote the analogous belief based on the stationary measures of queue lengths described in Proposition 3, that is ignoring the individual's prior on the calendar date. If the individual were certain that the queue had been operating long enough to be in the stationary regime she would form the posterior belief:

$$(14) \quad \bar{\mu}_n^0 := \frac{\mu y_n}{\mu y_n + (1 - \mu)w_n}.$$

For $n \leq \min(N, M)$, $\bar{\mu}_n^0$ depends on n only through y_n . Because $\bar{\mu}_n^0$ is an increasing function of y_n , the results on the form of the stationary measure in the previous section imply

that $\bar{\mu}_n^0$ is decreasing, constant, increasing in n for the same values of α as y_n is when $n \leq \min(N, M)$, i.e. for $\alpha > \alpha_N^*$, $\alpha = \alpha_N^*$ and $\alpha < \alpha_N^*$ respectively.

We now turn to the individual's private learning, or experimentation. As an individual waits in line, she observes whether those in front of her are served or not. As soon as service occurs, all agents currently in the queue learn that the server is in the good state and their posteriors jump to unity. We let μ_n^t denote the posterior of the individual at the n^{th} position ($n \leq N$) who has observed $t = 0, 1, 2, \dots$ unsuccessful service events. As before we will use $\bar{\mu}_n^t$ to denote the stationary analogue. From Bayes' rule this is:

$$(15) \quad \bar{\mu}_n^t := \frac{\bar{\mu}_n^0(1-\alpha)^t}{\bar{\mu}_n^0(1-\alpha)^t + 1 - \bar{\mu}_n^0}.$$

In Lemma 2 we show that the true posterior beliefs μ_n^t , for $t = 0, 1, 2, \dots$, can be made arbitrarily close to these posteriors if ν is chosen to be sufficiently small. The proof of this Lemma is given in Appendix A.5. The intuition for the proof is that the true posterior beliefs are an average of beliefs the agent would have formed if she knew she had arrived at a given calendar date. As $\nu \rightarrow 0$ this average gets closer to the time (or ergodic) average which uses the stationary measure. This is aided by the simple form that learning takes in this model: either posteriors jump to unity or they are revised downwards. Importantly this result is independent of q .

Lemma 2 *For any $M, \varepsilon > 0$ there exists a $\bar{\nu} > 0$, such that for all $\nu < \bar{\nu}$, $N < M$, $q \in [0, 1]$, $t \geq 0$:*

$$|\bar{\mu}_n^t - \mu_n^t| < \varepsilon.$$

We now compare agents' posterior beliefs along any given queue. Those ahead of the individual n^{th} in line may (for $\alpha > \alpha_N^*$) have been more optimistic than her when they joined the queue because they arrived at a shorter queue. However, they have been waiting in the queue for longer, and unless they have observed a service event, the waiting will have depressed their belief about the server state. We show that in any given queue in which no one has observed a service event, the most optimistic agent is the last in line. Lemma 3 establishes that an agent who is at the n^{th} position and has been waiting t periods is less optimistic than the agent at the $n+1^{\text{th}}$ position who has been waiting $t-1$ periods, if neither agent has observed service. The intuition for this result follows from the nesting of agents' information partitions. The individual behind n in the queue has observed strictly less than n has, so their beliefs about the state of service are an expectation of n 's beliefs. This expectation places positive weight on n knowing that the server is in the good state. That is, μ_{n+1}^t is an average of 1 and one point $\mu_n^{t+1} < 1$. Such an average must be above μ_n^{t+1} . 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. To be precise about our result, proved in Appendix A.6, we introduce a new piece of notation: let $\mu_n^{t\tau}$ be the posterior of an agent who is at the n^{th} position in the queue at calendar date τ and who has been queueing for t periods.

Lemma 3 *If $\mu_n^{t\tau} < 1$ then $\mu_{n+1}^{t-1\tau} > \mu_n^{t\tau}$ for all τ , n and $t > 0$.*

Our final lemma in this section describes some properties of the first in line's steady-state posterior, $\bar{\mu}_1^t$, as the parameters N and α of the strategy vary. (Below we will use the notation $\bar{\mu}_1^t(N)$ to make explicit the dependence of the beliefs on the parameter N .)

In Lemma 4(a), we describe some properties of the informational externality for the agent arriving at the queue in first position. For given values of M and N , there is a critical value $\bar{\alpha}$ such that if $\alpha > \bar{\alpha}$ finding herself first in line causes an agent to revise upwards her belief about the server state: $\mu_1^0(N) > \mu$, i.e. it is good news to be first in line. Whereas, for $\alpha < \bar{\alpha}$ being first in line is bad news and $\mu_1^0(N) < \mu$. The intuition for this result is that when α is large the stationary distribution in good states has a peak at $n = 1$, so on observing a short line the most likely explanation is that the state is good. Conversely when α is small, the stationary distribution in good states has peaks at $n = M + 1$ and on observing a short line the most likely explanation is that the server is in the bad state. Although the value of this threshold varies with N and M , we construct bounds on the threshold that are independent of these parameters.

In Lemma 4(b) we show that, as the strategy prescribes that the first in line experiment for longer (i.e. N increases) the probability of arriving at the queue at the first position when the server is good declines. An intuition is that as N increases there is total probability being spread over more states, so the probability of any one state falls.

Finally in Lemma 4(c), we show that although a higher N does affect the first in line's social learning, it still results in a reduction in the first in line's posterior after N unsuccessful service opportunities. As N increases there are many things to take account of: The probability 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. On 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. All these results are proved in Appendix A.7.

Lemma 4 *Suppose that $N > 1$ and $q = 1$, then:*

(a) *There exists a threshold value $\bar{\alpha} \in (0, 1)$ such that for each N*

$$\bar{\mu}_1^0(N) > \mu \quad \iff \quad \alpha > \bar{\alpha}.$$

Where $(3 - \sqrt{5})/2 \leq \bar{\alpha} \leq 2/3$ for all $M, N > 1$.

(b) *y_1 decreases as N increases for all $\alpha \in (0, 1)$.*

(c) *$\bar{\mu}_1^N(N)$ decreases in N for all $N > 1/\alpha$ and tends to zero as N tends to infinity.*

5.4 Equilibrium with Perfect Revelation

In this section we establish that for δ sufficiently large, the strategy $\sigma^*(q^*, N^*, M^*)$ constitutes a symmetric equilibrium for some $N^* < M^*$. That is, we complete the proof of Proposition 2. There are three conditions that need to be satisfied by an equilibrium strategy profile $\sigma^*(q^*, N^*, M^*)$.

First, because for $N < M$ the queue length reveals that the server is in the good state to the player arriving in M^{th} position, the equilibrium value M^* must equal \mathcal{M} , the longest

individually rational queue length defined in (2):

$$M^* = \mathcal{M} = \left\lfloor \frac{\ln(\delta w)}{\ln(2 - \delta)} \right\rfloor.$$

Second (Lemma 5), the equilibrium values of N and q determine the first in line's posterior μ_1^0 upon arriving at the queue at the first position and are an optimal policy for the first in line given that belief. This condition is summarized by the equation

$$N^* = \mathcal{N}(\mu_1^0(N^*), 1).$$

Third (Lemma 7), given the stationary measures of queue lengths generated by the strategy $\sigma^*(q^*, N^*, M^*)$, an individual arriving at the n^{th} position in the queue is prepared to herd on the first in line's actions. Reneging on the queue when the first in line reneges is clearly optimal. As once the first in line reneges, any player behind her becomes as pessimistic as her (adopts the posterior belief μ_1^N) and faces at least as much congestion as the first in line did. Therefore, it is sufficient to ensure that no individual in the queue wants to renege before the first in line's N^* periods of experimentation are completed.

The following lemma, proved in Appendix A.8, establishes the existence of an equilibrium number of periods N^* for the first in line to experiment. The intuition for this result comes from Lemma 4(c). There we show that as the first in line experiments for more periods, her posterior after unsuccessful experimentation must eventually converge to zero. This occurs even allowing for the effect which her increased experimentation has on the stationary distributions and on her resultant belief upon arriving at the queue. The continuity of the first in line's posterior μ_1^N in the strategy $\sigma^*(q, N, M)$ (again allowing for the effects on the stationary distributions), then ensures that there exists a posterior which hits the optimal exit threshold for the first in line.

Lemma 5 *There exists (q^*, N^*) , such that it is optimal for the first in line to wait N^* periods (if $q^* = 1$) or $N^*, N^* + 1$ periods (if $q^* < 1$) for service when all other agents use the strategy $\sigma^*(q^*, N^*, M)$ for any $M > N^*$.*

There are several points worth making about Lemma 5. The first is that the pair (q^*, N^*) is not necessarily unique for any given M . The possibility of multiple equilibria arises because increased experimentation by the first in line (changing the stationary measures in the two server states) results in an increased prior on the good state for the first in line. The increased prior then makes this increased experimentation optimal. Of course this process cannot continue indefinitely, by Lemma 4(c), so the set of possible equilibrium values of (N^*, q^*) is finite, but there is no clear monotonicity that ensures uniqueness.

Second, we can compare the equilibrium experimentation, (q^*, N^*) , with the single-agent optimum experimentation at the original prior μ . We know, from Lemma 4(a), that for all $N > 1$ and $\alpha > \bar{\alpha}$ the first in line's posterior at the equilibrium is above the prior μ . Thus optimality requires that she experiments more in this equilibrium of the game than in the corresponding single-agent decision problem. Similarly, for values of α below $\bar{\alpha}$, the first in line's posterior is less than her prior, so she must experiment less in the game than in the decision problem.

Corollary 1 *If m^* is given by $\mathcal{N}(\mu, 1)$, then: For all $\alpha > 2/3$ the equilibrium strategy $\sigma^*(q^*, N^*, M)$ with $N^* = \mathcal{N}(\mu_0^1, 1)$ satisfies $N^* \geq m^*$ and for all $\alpha < (3 - \sqrt{5})/2$ the equilibrium strategy satisfies $N^* \leq m^*$.*

Before proceeding we give an intermediate result on the limiting behaviour of N^* and M^* as $\delta \rightarrow 1$ for different values of α . Unambiguously, M^* tends to $+\infty$ for all $\alpha \in (0, 1)$. For $\alpha > 1/2$, as $\delta \rightarrow 1$, the willingness of the first in line to experiment also grows without bounds, even for high values of α (at which a good server should produce service quickly and the agents' posteriors quickly fall quite low if no service is observed), or for lower values of α ($1/2 < \alpha < \bar{\alpha}$) for which arriving first in line depresses an agent's belief relative to her prior.

For $\alpha < 1/2$, this last, "bad news" effect dominates, and $\lim_{M \rightarrow +\infty} y_1 = 0$ for $N > 1$ so that arriving at the first position in line makes an individual almost certain that the server is bad, and we have $N^* = 1$.¹⁴ For $\alpha = 1/2$, the effects balance out and N^* tends to some finite constant c that increases with the prior μ . We state our next Lemma with reference to an auxiliary problem described, together with the proof of the Lemma, in Appendix A.9.

Lemma 6 *For all $\mu \in (0, 1)$, for all $\alpha \in (0, 1)$, as $\delta \rightarrow 1$, N^* converges to*

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

The final step is to address our third equilibrium condition and show that no individual in the queue wishes to renege before the first in line's experimentation has elapsed. The next lemma provides a sufficient condition on the parameters of the model for this to be the case. The intuition for this result follows from the following simple argument. Observing the first in line's decision on whether or not to renege after N ($N + 1$ with probability $1 - q$) periods of experimentation reveals all the first in line's information to those behind her in the queue. So waiting for the result of the first in line's experimentation generates an informational benefit for later arrivals. The expected cost of acquiring this information is less for them than for the first in line, because even if they learn that the server is good later arrivals do not need to queue for as long as her in the bad server state. Their expected benefit is also lower, because later arrivals have to bear the cost of congestion parametrised by ψ . However, as the congestion becomes small ($\delta \rightarrow 1$) the discrepancy in benefits vanishes and therefore it ultimately becomes optimal to pay the reduced cost of observing the first in line.

Lemma 7 *Given $\mu \in (0, 1)$ and $\alpha \in (0, 1)$, there exists a $\underline{\delta}$ such that for all $\delta > \underline{\delta}$ and all $0 < \nu < \bar{\nu}(\delta)$, when all other agents use the strategy $\sigma^*(q^*N^*, M^*)$ given by Lemma 5 it is optimal for the n^{th} in line, $n = 2, \dots, M^* + 1$, to play this strategy.*

¹⁴In that case we also have $q^* < 1$. See the discussion in the Appendix A.9.

5.5 Other Equilibria

For the existence of an equilibrium with perfect revelation we required that δ be sufficiently large. Intuitively, agents must be sufficiently patient, or equivalently the congestion externality ψ they perceive must be sufficiently low, for them to accept taking up late positions in the queue. As δ decreases, individuals are more reluctant to take up late positions in the queue, and we can envisage an equilibrium in which no individual is willing to join the queue at position n or greater, but in which the first in line finds it optimal to experiment for $N > n$ periods. In this section we illustrate the existence of equilibria with imperfect revelation for intermediate¹⁵ values of δ .

At a strategy profile $\sigma^*(q, N, M)$ with $M \leq N$, the player joining the queue at the M^{th} position cannot learn, merely by observing the queue length, that the server is in the good state. Therefore, the condition $M^* = \mathcal{M}$ does not determine the equilibrium value of M , as it did in the equilibrium with perfect revelation. Instead, M^* is determined as follows: the individual arriving at the queue at the $M^* + 1$ st position must prefer balking, while the individual arriving at the M^{th} position must prefer joining the queue and waiting $N^* + 1 - M^*$ periods (with probability q , and $N^* + 2 - M^*$ periods with probability $1 - q$) to obtain the first-in-line's information.

For $n = 1, \dots, M$, the payoff from joining the line at n^{th} position and abiding by strategy $\sigma^*(q, N, M)$ is:

$$(16) \quad U^*(n) := q \left[\left(1 - \bar{\mu}_n^0 + \bar{\mu}_n^0 (1 - \alpha)^X \frac{z_n}{y_n} \right) \delta^X + \bar{\mu}_n^0 \left(1 - (1 - \alpha)^X \delta^X \frac{z_n}{y_n} \right) \psi^n \delta w \right] \\ + (1 - q) \left[\left(1 - \bar{\mu}_n^0 + \bar{\mu}_n^0 (1 - \alpha)^{X+1} \frac{z_n}{y_n} \right) \delta^{X+1} + \bar{\mu}_n^0 \left(1 - (1 - \alpha)^{X+1} \delta^{X+1} \frac{z_n}{y_n} \right) \psi^n \delta w \right],$$

where $X := N + 1 - n$ is the number of periods the n^{th} in line must wait before the first in line's N periods of experimentation are over, and where $\frac{z_n}{y_n} = (1 - \alpha)^{n-1} \frac{y_1}{y_n}$ is the likelihood which the player who arrives at the queue at the n^{th} position attributes to the first in line being uninformed for $n = 1, \dots, M$.

Within each set of square brackets, the first term is the payoff to the n^{th} in line if she eventually reneges on the queue, and the second term is her payoff if she is eventually served. The first term in round brackets is the probability that the player eventually reneges under strategy σ^* . For the first in line ($n = 1$) this equals the probability that the server is bad, or that it is good but produces N failures. For the n^{th} in line, $N = 2, \dots, M$, this equals the probability that the server is bad; or that the server is good yet produces $N + 1 - n$ failures *and the first in line is uninformed* (the individual arriving at the queue at the n^{th} position attributes probability z_n/y_n to that last even).

The individual arriving at the queue at the $M + 1$ st position could be the n^{th} arrival equiprobably for all $n \in \{M + 1, \dots, N\}$. Therefore, her expected payoff from joining the

¹⁵ Notice that for $\delta < (1 + \alpha(w - 1))^{-1}$, 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 therefore optional for any individual arriving at the queue to balk immediately and take the outside option.

queue at the $M + 1$ st position is

$$U_{M+1}^*(N, M) := \frac{q}{N - M} \sum_{n=M+1}^N U_{M+1}^*(n, X) + \frac{1 - q}{N + 1 - M} \sum_{n=M+1}^{N+1} U_{M+1}^*(n, X + 1),$$


where


$$U_{M+1}^*(n, X) := \left(1 - \bar{\mu}_{M+1}^0 + \bar{\mu}_{M+1}^0 (1 - \alpha)^X \frac{z_n}{y_{M+1}} \right) \delta^X + \bar{\mu}_{M+1}^0 \left(1 - (1 - \alpha)^X \delta^X \frac{z_n}{y_{M+1}} \right) \psi^{M+1} \delta w,$$

and where $X := N + 1 - n$ and $z_n = (1 - \alpha)^{n-1} y_1$ for $n = M + 1, \dots, N, N + 1$.

Under imperfect revelation, the equilibrium values q^* , N^* and M^* therefore satisfy the following conditions: (1) $N^* = \mathcal{N}(\bar{\mu}_1^0(N^*, M^*), 1)$, where $\bar{\mu}_1^0(N^*, M^*)$ depends on N^* and M^* via y_1 ; (2) $U^*(M^*) \geq 1$ and $U_{M^*+1}^* < 1$; (3) for $n = 2, \dots, M^*$, the player arriving n^{th} in line does not want to renege before the first in line's experimentation is completed.

Below we illustrate the equilibria of the queuing game as a function of δ for different values of α . Notice that there can be multiple equilibria for some pairs of parameters α, δ . For δ sufficiently high, an equilibrium with perfect revelation always exists. It ceases to exist as delta decreases, and instead an equilibrium with imperfect revelation exists. The transition is not sharp, and there can be parameter values at which both types of equilibria exists.





Equilibria in the queueing game for $\mu = 0.99$ and $\alpha = (0.7, 0.5, 0.3)$.

Finally, observe that the strategy σ^* we have studied here belongs to a broader class of strategies which focus on particular individuals at particular positions in the queue whose actions are informative. We will call these individuals “herding leaders”. The strategy of a herding leader is to pick a duration for which 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 herding followers is to focus on the closest herding leader ahead of them in the queue and to renege only when she does. So once in the queue, only a herding leader’s strategy depends on her private learning, whereas a herding followers’ strategies depend only on the publicly observed herding leaders’ actions. We could envisage the existence of equilibria with more than one herding leader, but we do not analyse it in this paper.

6 Social Memory

Let us define social memory to be the average time it takes to go from a state in which an individual arrives at the queue in first position to the next such state. An individual arriving first in line has no way of learning from the past experience of those who have been in the queue before her: the social memory is reset. By the standard results for positive recurrent Markov processes, the mean return time to the state in which an individual arrives at the queue in first position conditional on the server being good is given by $1/y_1$, the inverse of the stationary probability of that state (see for example, Brémaud (1999) p. 104). As this is something we have calculated (see A.21) we have the following result.

Corollary 2 *If the server is in the good state, the social memory is*

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

It is simple to see that as α approaches unity (ϕ goes to zero) the social memory vanishes, so the episodes in between social memory resets become very short. We can also show that these become arbitrarily large as α becomes small ($\phi \rightarrow \infty$).¹⁶ When this is the

¹⁶The greatest power of ϕ dominates the polynomial, this is $\phi^{M+1}(1 - \phi^N(\phi+1)^{-N})/(\phi-1)$. L’Hôpital’s rule shows that this tends to infinity.

case it must be that most of the time everyone in the queue knows that the server is in the good state

7 Conclusions and Further Work

The ingredients in our queueing model—individual learning, observational learning and payoff externalities—arise in many economic and social contexts. Consider, for example, firms that are engaged in R&D projects in closely 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 the number of firms that are competing in the area, the less lucrative 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.

Even though the nature of congestion in some of these contexts may be somewhat differently structured, similar issues as in our model arise, and we hope that the results derived here will be useful in analyzing these related problems. The crucial benefit of the queueing 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 queueing model offers to the more general question of experimentation with informational and payoff externalities is that strategy profiles in which individuals concentrate the social learning on certain focal individuals might result in such nesting of information and are more likely to constitute equilibria in more general settings.

References

- BANERJEE, A. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107, 797–817.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100, 992–1026.
- BOLTON, P., AND C. HARRIS (1999): "Strategic Experimentation," *Econometrica*, 67(2), 349–374.
- BRÉMAUD, P. (1999): *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer-Verlag, New York.
- BULOW, J. I., AND P. KLEMPERER (1994): "Rational Frenzies and Crashes," *The Journal of Political Economy*, 102(1), 1–23.
- CHAUDHRY, M. L., AND U. C. GUPTA (1996): "Performance Analysis of the Discrete-Time $GI/Geom/1/N$ Queue," *Journal of Applied Probability*, 25, 307–324.

- DEBO, L. G., C. PARLOUR, AND U. RAJAN (2012): “Signaling quality via queues,” *Management Science*, 58(5), 876–891.
- DEBO, L. G., U. RAJAN, AND S. VEERARAGHAVAN (2012): “Signaling by price in a congested environment,” *Chicago Booth Research Paper*, (12-13).
- EYSTER, E., A. GALEOTTI, N. KARTIK, AND M. RABIN (2013): “Congested Observational Learning,” .
- HASSIN, R. (1985): “On the Optimality of First Come Last Served Queues,” *Econometrica*, 53(1), 201–202.
- HASSIN, R., AND M. HAVIV (2003): *To Queue or not to Queue: Equilibrium Behavior in Queuing Systems*. Kluwer Academic Publishers, Dordrecht, The Netherlands.
- HERRERA, H., AND J. HÖRNER (2013): “Biased social learning,” *Games and Economic Behavior*.
- KELLER, G., S. RADY, AND M. CRIPPS (2005): “Strategic experimentation with exponential bandits,” *Econometrica*, 73(1), 39–68.
- MURTO, P., AND J. VÄLIMÄKI (2011): “Learning and information aggregation in an exit game,” *The Review of Economic Studies*, 78(4), 1426–1461.
- PERCUS, O. E., AND J. K. PERCUS (1990): “Elementary Properties of Clock-Regulated Queues,” *Siam Journal Applied Mathematics*, 50(4), 1166–1175.
- SMITH, L., AND P. SØRENSEN (2000): “Pathological Outcomes of Observational Learning,” *Econometrica*, 68(2), 371–398.
- (2009): “Biased Social Learning,” Discussion paper, University of Michigan.
- STRULOVICI, B. (2010): “Learning while voting: Determinants of collective experimentation,” *Econometrica*, 78(3), 933–971.
- THOMAS, C. (2013): “Strategic Experimentation with congestion,” Discussion paper, Mimeo, University of Texas at Austin.
- TOXVAERD, F. (2008): “Strategic merger waves: A theory of musical chairs,” *Journal of Economic Theory*, 140(1), 1–26.

A Appendix

A.1 Proof of Proposition 1

Proof: Taking a difference and substituting for ψ gives:

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

The term in braces is strictly decreasing in m and tends to negative infinity as $m \rightarrow \infty$. The function $U_n(\cdot, \mu_n^0)$ is, therefore, strictly quasi-concave in m and has a maximal value on $m \geq 0$. Thus, there is a solution to the problem $\max_{m \geq 0} U_n(m, \mu_n^0)$. 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 non-positive. This 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 (6).

After observing m periods of unsuccessful experimentation the individual forms the posterior belief

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

Using this expression and setting the braces in (A.1) equal to zero gives us the expression in (7) for $\underline{\mu}_n$, the n^{th} in line cutoff posterior. It is optimal for the n^{th} in line to experiment as long as $\mu_n^m \geq \underline{\mu}_n$ and to renege otherwise. \square

A.2 Proof of Proposition 3

Proof: I) Good server under perfect revelation ($N \leq M$): We will begin by considering the recursions which the stationary distribution of the queue lengths must satisfy when $N \leq M$.

Consider first the state in which the queue length is $n = 1$. It is possible to enter this state if there were previously r individuals in line and more than r service events occurred (probability α^r). It is also possible to enter state $n = 1$ if there were N or $N + 1$ individuals in line in the previous period and the first in line had never observed service, was not served and renege, 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 $n-1$ individuals in the line, or if $r - (n-1)$

individuals 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 his randomising 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 $r - N$ individuals 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 re-arranging 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

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

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 $y_M + y_{M+1}$, the probability that the latest agent arrives at the M^{th} position and joins the queue, or at the $M + 1^{\text{st}}$ position and balks. An agent arrives at the M^{th} 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 $M + 1^{\text{st}}$ position if the queue was of length M at the end of the last period and no service occurred:

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

Re-arranging 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 summarised below:

$$(A.2) \quad y_n = \begin{cases} \sum_{r=1}^M \alpha^r y_r + \alpha^M y_{M+1} + z_N(1-\alpha)(1-\alpha(1-q)), & n = 1; \\ (1-\alpha) \sum_{r=n-1}^M \alpha^{r-n+1} y_r + (1-\alpha)\alpha^{M-n+1} y_{M+1}, & 1 < n \leq N; \\ (1-\alpha) \sum_{r=N}^M \alpha^{r-N} y_r + (1-\alpha)\alpha^{M-N} y_{M+1} - q(1-\alpha)z_N, & n = N + 1; \\ (1-\alpha) \sum_{r=N+1}^M \alpha^{r-N-1} y_r + (1-\alpha)\alpha^{M-N-1} y_{M+1} - (1-q)(1-\alpha)^2 z_N, & n = N + 2; \\ (1-\alpha) \sum_{r=n-1}^M \alpha^{r-n+1} y_r + (1-\alpha)\alpha^{M-n+1} y_{M+1}, & N + 2 < n < M; \\ (1-\alpha)y_{M-1} + \alpha y_{M+1}, & n = M. \\ (1-\alpha)\alpha^{-1} y_M, & n = M + 1. \end{cases}$$

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

Before solving this system we will determine the value of z_N , 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 individual 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 individual is $(1 - \alpha)^{N-1} y_1$. Following the same argument for queue lengths $n \leq N$, we conclude that

$$(A.3) \quad z_n = (1 - \alpha)^{n-1} y_1, \quad n = 1, 2, \dots, N.$$

It is now clear that the system (A.2) is homogenous degree one.

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.2):

$$(A.4) \quad y_n = \begin{cases} \alpha(1 - \alpha)^{-1} y_2 + z_N(1 - \alpha)(1 - \alpha(1 - q)), & n = 1; \\ (1 - \alpha) y_{n-1} + \alpha y_{n+1}, & 1 < n < N; \\ (1 - \alpha) y_{N-1} + \alpha y_{N+1} + \alpha(1 - \alpha) q z_N, & n = N; \\ (1 - \alpha) y_N + \alpha y_{N+2} - (1 - \alpha) q z_N + \alpha(1 - \alpha)^2 (1 - q) z_N, & n = N + 1; \\ (1 - \alpha) y_{N+1} + \alpha y_{N+3} - (1 - \alpha)^2 (1 - q) z_N, & n = N + 2; \\ (1 - \alpha) y_{n-1} + \alpha y_{n+1}, & N + 2 < n < M + 1; \\ (1 - \alpha) \alpha^{-1} y_M, & n = M + 1. \end{cases}$$

We will now solve this difference equation. For $n = 1, 2, \dots, N$ 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.4). 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. (We treat the case where $\alpha = 1/2$ in Appendix A.3.)

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

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

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

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

Now let us turn to states $N + 2 < n \leq M + 1$. Taking the terminal condition given by the expression for y_M and y_{M+1} in (A.4) and substituting into the y_{M-1} equation gives $y_{M-1} = \left(\frac{\alpha}{1-\alpha}\right)^2 y_{M+1}$. Hence, $y_n = (\alpha/(1 - \alpha))^{M+1-n} y_{M+1}$. Or 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

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

We now substitute the value of z_N into the y_1 equation. A re-writing of (A.3) 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

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

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:

$$(A.5) \quad y_n = H \begin{cases} \phi^n - k_N \phi, & n = 1, 2, \dots, N; \\ \phi^n - k_N \left(\phi + \frac{q(1-\phi)}{\alpha(\phi+q)} \right), & n = N + 1; \\ \phi^n - k'_N \phi^{n-N-1}, & n = N + 2, \dots, M + 1; \end{cases}$$

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.5), are all non-negative. The terms y_{N+2}, \dots, y_{M+1} are all proportionate, so it is sufficient to check that y_1, \dots, y_{N+2} are non-negative. To address this question we will consider three separate cases.

Let α_N^* satisfy $(1 - \alpha + \alpha q)(2\alpha - 1) = (1 - \alpha)^{N+1}$. (Then, $\alpha_N^* < 1/2$ and $\alpha_N^* \rightarrow 1/2$ as $N \rightarrow \infty$.) Furthermore, $k_n < 0$ if $\alpha < \alpha_N^*$ and $k_n > 0$ if $\alpha > \alpha_N^*$. Also k_N is strictly decreasing when $\alpha > \alpha_N^*$ 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:

Case 3.1 ($1/2 < \alpha < 1$): Since $0 < k_N < 1$, to ensure $y_1 \geq 0$ we require $H \geq 0$. When $H > 0$ the terms y_1, \dots, y_{N+2} decrease (since $\phi < 1$), so it is sufficient to check that $y_{N+2} \geq 0$. This is the case since $\phi^{N+1} \geq k_N$.

Case 3.2 ($\alpha_N^* < \alpha < 1/2$): Since $k_N > 1$, from $y_1 \geq 0$ we must have $H \leq 0$. The terms y_1, \dots, y_N , therefore, decrease (since $\phi > 1$). It is sufficient to check that $y_N, y_{N+1}, y_{N+2} \geq 0$. The first two follow from $k_N > \phi^{N+1}$. To verify that $y_{N+2} \geq 0$ 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 3.3 ($\alpha < 1/2$): Since $k_N < 1$, from $y_1 \geq 0$ we must have $H \geq 0$. Since $k_N < 0$ all y_n are then positive.

The constant H must be chosen so that the y_n defined in (A.5) 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 re-write the above as (9) 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 $n = 1, \dots, M+1$. The uniqueness of this stationary distribution follows from the fact that the strategy described induces an irreducible Markov process on the states $n = 1, \dots, M$.

II) Good server under imperfect revelation ($M \leq N$): Now the queue never grows longer than length M , even if the first in line is still experimenting, because no other agent is willing to join a queue longer than M . The probability of arriving at the $M+1^{\text{st}}$ position (and then balking) depends on whether the first in line is informed or not. If the first in line is uninformed and there are M 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:

$$(A.6) \quad y_{M+1} = \frac{1 - \alpha}{\alpha} (y_M - z_M(1 - \alpha)^{N-M}(1 - \alpha + q\alpha)).$$

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

$$(A.7) \quad y_1 = z_1 = \sum_{r=1}^M \alpha^r y_r + \alpha^M y_{M+1} + z_1(1 - \alpha)^N(1 - \alpha + q\alpha).$$

For $M < N$, the probability of arriving at the n^{th} position satisfies the same recursion as for $N \leq M$:

$$(A.8) \quad y_n = (1 - \alpha) \sum_{r=n-1}^M \alpha^{r-n+1} y_r + (1 - \alpha) \alpha^{M-n+1} y_{M+1}, \quad n = 2, \dots, M;$$

and the probability of arriving at the n^{th} position and the first in line being uninformed is

$$(A.9) \quad z_n = (1 - \alpha)^{n-1} z_1, \quad n = 2, \dots, M.$$

The recursion A.8 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. (We treat the case where $\alpha = 1/2$ in Appendix A.3.) Rewriting the initial condition A.7 by substituting A.8 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:

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

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.6 (where we use A.9 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

$$(A.10) \quad 1 = \frac{z_1}{1 - k_N} \left(\sum_{n=1}^{M-1} (\phi^{n-1} - k_N) + \frac{\phi^M - k_N}{1 - \alpha} \right),$$

which simplifies to

$$(A.11) \quad 1 = \frac{z_1}{1 - k_N} \left(\frac{1 - \phi^{M+1}}{1 - \phi} - k_N(M + \phi^{-1}) \right).$$

This determines the last part of the solution.

We now verify that all the y_n are non-negative. For $2 \leq n \leq N + 1$, 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 $M \leq N$ 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 $z_1 \geq 0$.

From (A.11) we have that $z_1 > 0$ for $\alpha < \alpha_N^*$, because $k_N < 0$. In (A.10) 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 $z_1 \geq 0$ also for $\alpha > \alpha_N^*$. Hence we have derived a legitimate stationary measure when $M \leq N$.

III) Bad server:

We conclude by deriving the stationary distribution conditional on the server being in the bad state. For $M \leq N$ the transition equations are: $w_1 = \dots = w_N$ and $w_{N+1} = (1 - q)w_N$. For $N < M$ they are: $w_1 = \dots = w_M$ and $w_{M+1} = (N - M + 1 - q)w_M$. In each case the result follows from the requirement that the probabilities sum to 1. \square

A.3 Stationary distribution for $\alpha = 1/2$.

Lemma 8 *Let $\alpha = 1/2$. For $N < M$, the stationary distribution of queue lengths is*

$$y_n = \begin{cases} \frac{2(2^{N+1} - (n-1)(1+q))}{(M+1)2^{N+2} - (1+q)N(2M-N+1) - 4(M-N+q)}, & n \leq N, \\ \frac{2(2^{N+1} - N(1+q)) - 4q}{(M+1)2^{N+2} - (1+q)N(2M-N+1) - 4(M-N+q)}, & n = N + 1, \\ \frac{2(2^{N+1} - N(1+q)) - 4}{(M+1)2^{N+2} - (1+q)N(2M-N+1) - 4(M-N+q)}, & n \geq N + 2. \end{cases}$$

For $N \geq M$, the stationary distribution of queue lengths is

$$y_n = \begin{cases} \frac{2(2^{N+1} - (n-1)(1+q))}{(M+1)2^{N+2} - (1+q)((M+1)M+2)}, & n \leq M, \\ \frac{2(2^{N+1} - (M+1)(1+q))}{(M+1)2^{N+2} - (1+q)((M+1)M+2)}, & n = M + 1. \end{cases}$$

Proof: We now derive the stationary distribution of queue lengths $n = 1, \dots, M+1$ for the case where $N < M$, by solving the system of difference equations in (A.4) for the case where $\alpha = 1/2$. For $n = 1, 2, \dots, N$, y_n solves the difference equation $0 = (1-\alpha)y_{n-1} - y_n + \alpha y_{n+1}$, whose characteristic polynomial, $(x-1)(x-(1-\alpha)/\alpha)$, admits a unique root when $\alpha = 1/2$. We therefore obtain the general solution:

$$y_n = K + nH.$$

Imposing the initial condition, given by the expression for y_1 in (A.4), on this equation, we solve for H and obtain:

$$y_n = K - n z_N \frac{1}{4}(1+q), \quad n = 1, 2, \dots, N.$$

Substituting into the expressions for y_N , y_{N+1} and y_{N+2} in (A.4) respectively, we obtain:

$$\begin{aligned} y_{N+1} &= K - (N+1)z_N \frac{1}{4}(1+q) - z_N \frac{1}{2}q, \\ y_{N+2} &= K - (N+2)z_N \frac{1}{4}(1+q) - z_N \frac{1}{4}(1-q), \\ y_{N+3} &= K - (N+3)z_N \frac{1}{4}(1+q) + z_N \frac{1}{2}q. \end{aligned}$$

The terminal condition, given by the expression for y_{M+1} in (A.4), gives $y_M = y_{M+1}$, and from the expression for y_n when $N+2 < n < M+1$ in (A.4) we obtain that:

$$y_{M+1} = y_M = \dots = y_{N+3}.$$

Substituting the expression for y_1 into $z_N = (1-\alpha)^{N-1}y_1$ gives:

$$z_N = \zeta K, \quad \zeta := \frac{4}{2^{N+1} + 1 + q}.$$

Imposing that the y_n sum to unity:

$$1 = \sum_{n=1}^N y_n + y_{N+1} + y_{N+2} + \sum_{n=N+3}^{M+1} y_{N+3},$$

and solving for K we obtain:

$$K^{-1} = M + 1 - \zeta \frac{1}{8}(N+3)(2M-N)(1+q) + \zeta \frac{1}{2}(M-N-2)q - \zeta \frac{1}{4}(1-q).$$

The resulting stationary distribution of queue lengths when $N < M$ is described in the above lemma. (The case $N \geq M$ can be analysed in a similar fashion.) \square

A.4 Proof of Lemma 1

Proof: Let $\pi_{ss'}$ denote the probability of moving from state $s \in S$ to state $s' \in S$ under the Markov process followed by the queue in the good state. Also, let $\pi_s = \{\pi_{ss'}\}_{s' \in S} \in \Delta(S)$ denote the probability distribution of tomorrow's state $s' \in S$ conditional on today's state being s . Finally, let s^* denote the queue state where there is one uninformed individual in the queue at the end of the period.

We can bound the distance between two distributions π_s and π_r in the following way

$$\sum_{s' \in S} |\pi_{ss'} - \pi_{rs'}| \leq |\pi_{ss^*} - \pi_{rs^*}| + (1 - \pi_{ss^*}) + (1 - \pi_{rs^*}), \quad \forall s, r \in S.$$

(This upper bound follows by separating out the $s' = s^*$ term and then realizing that the remaining terms would be maximised if the support of π_s and π_r only had the point s^* in common.) Without loss of generality, suppose that $\pi_{ss^*} \geq \pi_{rs^*}$. If this is so, then substituting $|\pi_{ss^*} - \pi_{rs^*}| = \pi_{ss^*} - \pi_{rs^*}$ we have

$$\|\pi_s - \pi_r\| := \frac{1}{2} \sum_{s' \in S} |\pi_{ss'} - \pi_{rs'}| \leq 1 - \pi_{rs^*} \leq 1 - \alpha^M, \quad \forall s, r \in S.$$

The final inequality above follows from the construction of the service process: $\pi_{ss^*} \geq \alpha^M$ for all $s \in S$. The extremes of the above chain of inequalities imply that the Dobrushin Coefficient of this process is less than $1 - \alpha^M$. The Lemma then follows by Theorem 7.2, p.237, of Brémaud (1999). \square

A.5 Proof of Lemma 2

Proof: Individuals have the prior μ that the server is in the good state and the prior $\nu(1 - \nu)^\tau$ that they have arrived in the system at calendar date τ . Let y_n^τ (respectively w_n^τ) denote the probability that the individual arriving at calendar date τ finds herself at the n^{th} position in line, conditional on individuals using the postulated queueing strategy and the server being good (respectively bad). She would then attach probabilities

$$\begin{aligned} \mu \sum_{\tau=0}^{\infty} \nu(1 - \nu)^\tau y_n^\tau &:= \mu \beta_n^1 \\ (1 - \mu) \sum_{\tau=0}^{\infty} \nu(1 - \nu)^\tau w_n^\tau &:= (1 - \mu) \beta_n^2 \end{aligned}$$

to the server state being good or bad.

Recall that y_n denotes the stationary probability of queue length n in the good state. The following calculation shows that $|\beta_n^1 - y_n| \rightarrow 0$ as $\nu \rightarrow 0$.

$$|y_n - \beta_n^1| \leq \sum_{\tau=0}^{\infty} \nu(1 - \nu)^\tau |y_n - y_n^\tau|$$

$$\begin{aligned}
&\leq \sum_{\tau=0}^{\infty} \nu(1-\nu)^{\tau}(1-\alpha^M)^{\tau} \\
&= \frac{\nu}{(1-(1-\nu)(1-\alpha^M))} \rightarrow 0 \quad \text{as } \nu \rightarrow 0.
\end{aligned}$$

Where the second inequality follows from Lemma 1. (Note the rate of convergence here is independent of q .)

Recall that w_n denotes the stationary probability of queue length n in the bad state. We now prove that $|w_n - \beta_n^2| \rightarrow 0$ as $\nu \rightarrow 0$ at a rate that is independent of q . The sum $\sum_{\tau=t}^{t+N} w_n^{\tau}$ is the expected number of times the queue has length n in over the periods $t, \dots, t+N$. In $N+1$ consecutive periods, state $n < N$ must be visited at least once and can be visited twice if the initial state is n and state $N+1$ was not visited. This gives the two equalities:

$$(A.12) \quad \sum_{\tau=t}^{t+N} w_n^{\tau} = 1 + qw_n^t, \quad \sum_{\tau=t+1}^{t+N} w_n^{\tau} = 1 - (1-q)w_{n+1}^{t+1};$$

(the second inequality uses the fact that $w_n^t = w_{n+1}^{t+1}$.) We can use these to re-write β_n^2 . Let $S_n^{(N+1)\tau} := \frac{\nu}{1-(1-\nu)^{N+1}} \sum_{t=0}^N (1-\nu)^t w_n^{(N+1)\tau+t}$ then

$$\begin{aligned}
\beta_n^2 &= (1-(1-\nu)^{N+1}) \sum_{\tau=0}^{\infty} (1-\nu)^{(N+1)\tau} S_n^{(N+1)\tau} \\
&= (1-(1-\nu)^{N+1}) \sum_{\tau=0}^{\infty} (1-\nu)^{(N+1)\tau} \left(\frac{\sum_{t=0}^N w_n^{(N+1)\tau+t}}{N+1} + \left(S_n^{(N+1)\tau} - \frac{\sum_{t=0}^N w_n^{(N+1)\tau+t}}{N+1} \right) \right)
\end{aligned}$$

Now observe that $S_n^{(N+1)\tau} \rightarrow \frac{1}{N+1} \sum_{t=0}^N w_n^{(N+1)\tau+t}$ as $\nu \rightarrow 0$ and the order of convergence is $o(\nu)$. We then can substitute from (A.12) to get

$$\begin{aligned}
\beta_n^2 &= (1-(1-\nu)^{N+1}) \sum_{\tau=0}^{\infty} (1-\nu)^{(N+1)\tau} \frac{1 + qw_n^{(N+1)\tau}}{N+1} + o(\nu) \\
(A.13) \quad &= \frac{1}{N+1} + \frac{q}{N+1} (1-(1-\nu)^{N+1}) \sum_{\tau=0}^{\infty} (1-\nu)^{(N+1)\tau} w_n^{(N+1)\tau} + o(\nu)
\end{aligned}$$

By taking blocks of length N and making a different substitution from (A.12) we can also get

$$(A.14) \quad \beta_n^2 = \frac{1}{N} - \frac{1-q}{N} (1-(1-\nu)^N) \sum_{\tau=0}^{\infty} (1-\nu)^{N\tau} w_{n+1}^{N\tau} + o(\nu)$$

When $q = 0$ or $q = 1$, (A.13) and (A.14) are sufficient to prove that $|\beta_n^2 - w_n| \rightarrow 0$ as $\nu \rightarrow 0$. However, for $q \in (0, 1)$ more is required.

Now we apply similar reasoning to Lemma 1 for the stochastic process followed by the queues in the bad state. Consider a queue starting in state n and another queue starting

in state $n' > n$. In $(N+1)(n'-n)$ periods the queue starting in state n will be in state n' if it never visits state $N+1$. In the same number of periods the queue starting in state n' will return to state n' if it always visits state $N+1$. The first of these histories occurs with probability $q^{n'-n}$ the second occurs with probability $(1-q)^{n'-n}$. Thus after $(N+1)(n'-n)$ periods the initial states n and n' will be in the same state with at least probability $\underline{q}^{n'-n}$; where $\underline{q} := \min\{q, 1-q\}$. Hence after $N(N+1)$ periods there is at least probability \underline{q}^N that any two initial states result in the same current state. By Dobrushin's result we then have that

$$(A.15) \quad \sum_{n=1}^{N+1} |w_n^{N(N+1)t} - w_n| < (1 - \underline{q}^{N(N+1)})^t.$$

Finally, we consider $|\beta_n^2 - w_n|$. We begin by doing the case where $q > 1/2$ so $\underline{q} = 1 - q$. A substitution from (A.14) and $w_n = (N+1-q)^{-1}$ gives

$$\begin{aligned} |\beta_n^2 - w_n| &\leq \frac{1-q}{N+1-q} (1 - (1-\nu)^N) \sum_{\tau=0}^{\infty} (1-\nu)^{N\tau} \left| \frac{1}{N+1-q} - w_n^{N\tau} \right| + o(\nu) \\ &\leq \frac{1-q}{N+1-q} (1 - (1-\nu)^N) \sum_{\tau=0}^{\infty} (1-\nu)^{N\tau} K (1 - \underline{q}^{N(N+1)})^{\tau/(N+1)} + o(\nu) \\ &\leq \frac{K}{N+1-q} (1 - (1-\nu)^N) \sum_{\tau=0}^{\infty} (1-\nu)^{N\tau} (1-q) e^{-(1-q)\tau/(N+1)} + o(\nu) \\ &\leq \frac{K(N+1)}{e(N+1-q)} (1 - (1-\nu)^N) \sum_{\tau=1}^{\infty} \frac{(1-\nu)^{N\tau}}{\tau} + o(\nu) \\ &= -\frac{K(N+1)}{e(N+1-q)} (1 - (1-\nu)^N) \log(1 - (1-\nu)^N) + o(\nu) \end{aligned}$$

The second inequality here substitutes from (A.15) and introduces the constant K to accommodate the fact that (A.15) applies every $N(N+1)$ periods but we wish to bound every $N+1$ periods. The third inequality uses the fact that $1-x \leq e^{-x}$. The fourth inequality follows as $xe^{-x} \leq e^{-1}$ implies $(1-q)e^{-(1-q)\tau/(N+1)} \leq (N+1)/(\tau e)$ for $\tau > 0$ (the $\tau=0$ term can be included in the $o(\nu)$ factor). The final equality evaluates the sum $G(x) := \sum_{n=1}^{\infty} x^n/n$ by observing that $G'(x) = 1/(1-x)$. Hence we have that $|\beta_n^2 - w_n| \rightarrow 0$ as $\nu \rightarrow 0$ independently of q .

The final step in this proof is to apply the convergence results. The individual's true posterior upon arriving at a queue at the n^{th} position satisfies

$$\mu_n^0 = \frac{\mu\beta_n^1}{\mu\beta_n^1 + (1-\mu)\beta_n^2}.$$

We now compare this posterior to $\bar{\mu}_n^0$, the posterior based on the stationary distributions:

$$|\bar{\mu}_n^0 - \mu_n^0| = \left| \frac{\mu y_n}{\mu y_n + (1-\mu)w_n} - \frac{\mu\beta_n^1}{\mu\beta_n^1 + (1-\mu)\beta_n^2} \right|$$

$$\begin{aligned}
&\leq \left| \frac{\mu y_n}{\mu y_n + (1 - \mu)w_n} - \frac{\mu \beta_n^1}{\mu \beta_n^1 + (1 - \mu)w_n} \right| + \left| \frac{\mu \beta_n^1}{\mu \beta_n^1 + (1 - \mu)w_n} - \frac{\mu \beta_n^1}{\mu \beta_n^1 + (1 - \mu)\beta_n^2} \right| \\
&= \frac{\mu(1 - \mu)w_n |y_n - \beta_n^1|}{(\mu y_n + (1 - \mu)w_n)(\mu \beta_n^1 + (1 - \mu)w_n)} + \frac{\mu \beta_n^1 (1 - \mu) |w_n - \beta_n^2|}{(\mu \beta_n^1 + (1 - \mu)w_n)(\mu \beta_n^1 + (1 - \mu)\beta_n^2)} \\
&\leq \frac{\mu |y_n - \beta_n^1|}{(\mu \beta_n^1 + (1 - \mu)w_n)} + \frac{(1 - \mu) |w_n - \beta_n^2|}{(\mu \beta_n^1 + (1 - \mu)w_n)} \\
&\leq \frac{1}{\beta_n^1} |y_n - \beta_n^1| + \frac{1}{w_n} |w_n - \beta_n^2| \\
&\rightarrow 0 \text{ as } \nu \rightarrow 0.
\end{aligned}$$

To perform this calculation for individuals who have been in the system for t periods, i.e. for $|\mu_n^t - \bar{\mu}_n^t|$, it is necessary to account for an individual's learning from observing the behaviour of others and service events. Given the postulated strategies, an individual either depresses her belief because no service occurs, or become certain that the server is good (when service occurs or the first in line does not renege after N unsuccessful service events). When the individual is certain that the server is good, $\mu_n^t = 1$, and then also $\bar{\mu}_n^t = 1$ and the bound holds. When the individual has observed no service for t periods it is necessary to multiply β_1 by the factor $(1 - \alpha)^t$ to determine the updated posterior μ_n^t :

$$(A.16) \quad \mu_n^t = \frac{\mu \beta_1 (1 - \alpha)^t}{\mu \beta_1 (1 - \alpha)^t + (1 - \mu) \beta_2}.$$

The same factor multiplies μy_n in $\bar{\mu}_n^t$ (see (15)). Hence the bound above applies to $|\mu_n^t - \bar{\mu}_n^t|$ for all t . As the bound is a continuous function of ν the result follows. \square

A.6 Proof of Lemma 3

Proof: An individual's beliefs at calendar date τ about the state of service are an expectation: $\mu_n^{t\tau} = E(\mathbb{1}_{\text{good}} | h_{\tau-t}^\tau)$, where $\mathbb{1}_{\text{good}}$ is the indicator function for the event that the server state is good and $h_{\tau-t}^\tau$ describes the t periods of history that the agent who is in n^{th} position at date τ has observed if she has been queueing for t periods. (The history $h_{\tau-t}^\tau$ must be consistent with the n^{th} agent arriving at date $\tau - t$ and still being in line at date τ .) Notice that the $n + 1^{\text{st}}$ agent in line at date τ has observed strictly less information than the n^{th} in line (the history $h_{\tau-t}^\tau$ observed by the n^{th} in line includes the entire history $h_{\tau-t+1}^\tau$ observed by the $n + 1^{\text{st}}$ in line plus what the n^{th} in line observed in the period before the $n + 1^{\text{st}}$ arrived). By the nesting of the information sets we have

$$\mu_{n+1}^{t-1\tau} = E(\mathbb{1}_{\text{good}} | h_{\tau-t+1}^\tau) = E(E(\mathbb{1}_{\text{good}} | h_{\tau-t}^\tau) | h_{\tau-t+1}^\tau) = E(\mu_n^{t\tau} | h_{\tau-t+1}^\tau),$$

for any history consistent with the n^{th} agent arriving at date $\tau - t$ and still being present at date τ .

For $n \leq N + 1$, the variable $\mu_n^{t\tau}$ takes only two values: unity and $\mu_n^t < 1$ defined in (A.16). (It takes this value if the n^{th} in line has learnt nothing from others and has revised

downward her beliefs as a result of waiting for service.) Thus if we re-write the extremes above we have

$$\mu_{n+1}^{t-1\tau} = 1\pi(h_{\tau-t+1}^\tau) + (1 - \pi(h_{\tau-t+1}^\tau))\mu_n^t,$$

where $\pi(h_{\tau-t+1}^\tau)$ is the probability the $n+1^{\text{st}}$ in line attaches to the n^{th} in line being certain that the server is good. If $\mu_n^{t\tau} < 1$, then $h_{\tau-t}^\tau$ does not contain a service event. Therefore (i) by the nesting of information sets, neither does $h_{\tau-t+1}^\tau$ and $\pi(h_{\tau-t+1}^\tau) \in (0, 1)$, and (ii) $\mu_n^{t\tau} = \mu_n^t$. Then a substitution and a rearranging of the above gives:

$$\mu_{n+1}^{t-1\tau} - \mu_n^{t\tau} = \pi(h_{\tau-t+1}^\tau)(1 - \mu_n^t) > 0.$$

which proves the result. \square

A.7 Proof of Lemma 4

Proof: *Part (a):*

Assume that $M > N > 1$ and $q = 1$. From (14) we have that $\bar{\mu}_1^0 < \mu$ if and only if $Ny_1 < 1$. A substitution from (9) and (8) gives

$$\begin{aligned} \frac{1}{Ny_1} &= 1 + \frac{\frac{1-\phi^N}{1-\phi} - N + (1 - \frac{k_N}{\phi^{N+1}}) \frac{\phi^N - \phi^{M+1}}{1-\phi}}{N(1 - k_N)} \\ &= 1 + \frac{\sum_{i=0}^{N-1} (\phi^i - 1) + (1 - \frac{k_N}{\phi^{N+1}}) \frac{\phi^N - \phi^{M+1}}{1-\phi}}{N(1 - k_N)} \\ \text{(A.17)} \quad &= 1 + \frac{\frac{\phi^N - \phi^{M+1}}{1-\phi} \left(1 - \frac{1-\phi^{N+1}}{(1-\phi)(1+\phi)^N}\right) - \sum_{i=0}^{N-1} \frac{1-\phi^i}{1-\phi} \left(1 - \phi + \phi \left(\frac{\phi}{1+\phi}\right)^N\right)}{N} \end{aligned}$$

(To get the final line we substitute $k_N = \phi^{N+1}/(\phi^{N+1} + (1 - \phi)(1 + \phi)^N)$, when $q = 1$.) Notice that the first term in the numerator is positive for all $\phi > 0$, because: $M > N$ and $(1 + \phi)^N > \sum_{i=0}^N \phi^i$. Thus a necessary and sufficient condition for $Ny_1 > 1$ is that

$$\text{(A.18)} \quad \left(\frac{\sum_{i=0}^{N-1} \frac{1-\phi^i}{1-\phi}}{\frac{\phi^N - \phi^{M+1}}{1-\phi}} \right) \frac{1 - \phi + \frac{\phi^{N+1}}{(1+\phi)^N}}{1 - \frac{1-\phi^{N+1}}{(1-\phi)(1+\phi)^N}} > 1.$$

We will show that the left-hand side of (A.18) is decreasing in ϕ until it becomes negative and then stays negative for all larger ϕ . Therefore, there is a threshold value of ϕ such that (A.18) holds iff ϕ is below the threshold.

First consider the quotient in parentheses in (A.18). This can be written as

$$\frac{\sum_{i=0}^{N-1} \sum_{j=0}^{i-1} \frac{\phi^j}{\phi^N}}{\sum_{i=0}^{M-N} \phi^i}.$$

The denominator is increasing in ϕ and the numerator is decreasing in ϕ , so the term in parentheses in (A.18) decreases in ϕ for all $\phi > 0$.

Now suppose $\phi < 1$ and consider the second fraction in (A.18). The numerator equals $1 - \sum_{i=1}^N [\phi/(1+\phi)]^i > 0$, which decreases in ϕ . The denominator of this fraction has the derivative (in ϕ) equal to

$$\frac{N}{(1-\phi)(1+\phi)^N} \left[\frac{1+\phi^N}{1+\phi} - N^{-1} \frac{1-\phi^N}{1-\phi} \right].$$

Notice that $(1+\phi^N)/(1+\phi) \geq (1+\phi^{N-1})/2$ (with equality for $N = 1$ and strict inequality for $N \geq 2$) and $(1+\phi^{N-1})/2 \geq \sum_{i=0}^{N-1} \phi^i/N$ (with equality for $N = 1, 2$ and strict inequality for $N \geq 3$). Therefore, the difference above is non-negative and this derivative is non-negative when $\phi < 1$. Hence we have shown that the second fraction in (A.18) decreases when $\phi < 1$.

Now suppose $\phi > 1$. We write the second fraction in (A.18) as

$$(A.19) \quad \frac{(1-\phi)(1+\phi)^N + \phi^{N+1}}{(1+\phi)^N - \frac{1-\phi^{N+1}}{1-\phi}}.$$

The denominator of (A.19) is an n^{th} order polynomial in ϕ with positive coefficients so it is increasing in ϕ . The numerator of (A.19) has a derivative in ϕ that equals

$$-(1+\phi)^{N-1} \left\{ (N+1) \frac{\phi}{1+\phi} \frac{1 - \left(\frac{\phi}{1+\phi}\right)^{N-1}}{1 - \frac{\phi}{1+\phi}} + 1 - N \right\}$$

The term in braces increases in ϕ , thus it is smallest when $\phi = 1$. Evaluating these braces at $\phi = 1$ gives $2(1 - (N+1)2^{-N})$ which is positive for all $N > 1$. Thus this derivative is strictly negative. It follows that (A.19) decreases when $\phi > 1$ until the numerator becomes negative at which point (A.19) remains negative for all greater ϕ .

When $M \leq N$ substitutions from (11) and (10) give

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

When $\phi < 1$ (and $k_N < 1$) the top of the fraction is negative, so $y_1 > 1/(M+1)$ and $y_1 \geq 1/N$ unless $M = N$. Thus for all $\phi < 1$ we have that $\bar{\mu}_1 > \mu$. When $\phi > 1$ and $k_N > 1$ the top of the fraction is positive and the bottom is negative so still $y_1 > 1/(M+1)$. When $\phi > 1$ and $k_N < 0$ then a substitution for k_N gives

$$(A.20) \quad \frac{1}{y_1} = M+1 - \left(\frac{\phi}{1+\phi}\right)^N - \sum_{i=0}^M \frac{\phi^i - 1}{\phi - 1} \left(1 - \phi + \phi \left(\frac{\phi}{1+\phi}\right)^N\right)$$

Differentiation with respect to ϕ (and abbreviating the summation to Σ and the final parenthesis to A) gives

$$\frac{-N}{\phi^2} \left(\frac{\phi}{1+\phi}\right)^{N+1} + \Sigma - \Sigma \left(\frac{\phi}{1+\phi}\right)^N + \Sigma \frac{N}{\phi} \left(\frac{\phi}{1+\phi}\right)^{N+1} - \frac{\partial \Sigma}{\partial \phi} A$$

As A is negative (when $k_N < 0$) we have a lower bound on this derivative

$$\Sigma \left(1 - \left(\frac{\phi}{1+\phi} \right)^N \right) + \frac{N}{\phi} \left(\frac{\phi}{1+\phi} \right)^{N+1} \left[\Sigma - \frac{1}{\phi} \right] > 0, \quad \phi > 1.$$

Thus $1/[(M+1)y_1]$ increases in ϕ when $k_N < 0$ which is what we need to show.

The lower bound on $\bar{\alpha}$ follows from observing that the left of (A.18) is negative iff $1 - \phi + \phi(\phi/(1+\phi))^N < 0$. This is decreasing in N so tightest when $N = 2$, giving the inequality $1 + \phi < \phi^2$. The upper bound on $\bar{\alpha}$ follows from the observing that (when $\alpha > 1/2$) the unbracketed term in (A.18) is bounded above by $1 - \phi$ hence when $\alpha > 1/2$ a sufficient condition for (A.18) is

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

Letting $M \rightarrow \infty$ and setting $N = 2$ gives the sufficient condition $\phi < 1/2$. The upper bound then follows.

Part (b) Assume that $M > N > 1$. From (A.17) we get

$$\begin{aligned} \frac{1}{y_1} &= N + \frac{\phi^N - \phi^{M+1}}{1 - \phi} \left(1 - \frac{1 - \phi^{N+1}}{(1 - \phi)(1 + \phi)^N} \right) - \sum_{i=0}^{N-1} \frac{1 - \phi^i}{1 - \phi} \left(1 - \phi + \phi \left(\frac{\phi}{1 + \phi} \right)^N \right) \\ &= \sum_{i=0}^{N-1} \phi^i + \frac{\phi^N - \phi^{M+1}}{1 - \phi} \left(1 - \frac{1 - \phi^{N+1}}{(1 - \phi)(1 + \phi)^N} \right) - \phi \left(\frac{\phi}{1 + \phi} \right)^N \sum_{i=0}^{N-1} \frac{1 - \phi^i}{1 - \phi} \\ &= \sum_{i=0}^M \phi^i - \frac{\phi^N - \phi^{M+1}}{1 - \phi} \frac{1 - \phi^{N+1}}{(1 - \phi)(1 + \phi)^N} - \phi \left(\frac{\phi}{1 + \phi} \right)^N \sum_{i=0}^{N-1} \frac{1 - \phi^i}{1 - \phi} \\ &= \frac{1 - \phi^{M+1}}{1 - \phi} - \left(\frac{\phi}{1 + \phi} \right)^N \left\{ \phi \sum_{i=0}^{N-1} \frac{1 - \phi^i}{1 - \phi} + \frac{1 - \phi^{M-N+1}}{1 - \phi} \frac{1 - \phi^{N+1}}{1 - \phi} \right\} \end{aligned}$$

We now focus on the term in braces this equals

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

Hence

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

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

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

Then

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

(A substitution gives the sign.) When $M \leq N$ we have from (A.20)

$$(A.22) \quad \frac{1}{y_1} = \sum_{i=0}^M \phi^i - \left(\frac{\phi}{1+\phi}\right)^N \left[1 + \phi \sum_{i=0}^M \frac{\phi^i - 1}{\phi - 1}\right].$$

This implies y_1 decreases as N increases.

Part (c): To show that $\bar{\mu}_1^N$ to decreases in N it is sufficient to show that $N(1-\alpha)^N$ decreases in N as y_1 decreases in N from part (2). But $N(1-\alpha)^N$ decreases in N for all $N > 1/\alpha$. Finally, $(1-\alpha)^N y_1$ will converge to zero as N increases because $y_1 \leq 1$. \square

A.8 Proof of Lemma 5

Proof: In Proposition 1, (6) defines a map from priors $\mu \in (0, 1)$ to the optimal waiting time, $m^* \in \mathbb{Z}_+$, for the first in line

$$m^*(\mu) := \left\lceil (\log(1-\alpha))^{-1} \log \left(\frac{1-\mu}{\mu} \frac{\psi(1-\delta)}{\alpha(\psi\delta w - 1)} \right) \right\rceil_+.$$

This is an increasing step function from $(0, 1)$ to \mathbb{R}_+ that jumps upwards and converges to infinity as $\mu \rightarrow 1$. Define $m^\dagger(\mu)$ to be the staircase correspondence, from $(0, 1)$ to \mathbb{R}_+ , that includes the intervals where m^* jumps and everywhere else equals m^* . Finally, let f denote the inverse of m^\dagger , that is $f(x) := \{\mu : x \in m^\dagger(\mu)\}$. One can interpret $f(x)$ as the set of priors for which the optimal time to wait is x . (Waiting a non-integer time indicates that both the of the two nearest integers are optimal.) The correspondence f increases in x and $f(0)$ is the interval of priors $[0, \underline{\mu}]$ and as x tends to infinity $f(x) \rightarrow 1$. Where $\underline{\mu}$ is the largest prior for which it is optimal for the first in line to wait zero periods

$$\underline{\mu} := \frac{1-\delta}{\delta\alpha(w-1)}.$$

We now define a second function that maps \mathbb{R}_+ to $[0, 1]$. For $x \in \mathbb{R}_+$ consider the strategy σ that sets $N(x) = \lfloor x \rfloor$ and $q(x) = x - \lfloor x \rfloor$ (where $\lfloor x \rfloor$ is the greatest integer less than or equal to x) and otherwise satisfies the properties of Definition 1. The strategy $(N(x), q(x))$ determines a good-state stationary measure, $y(x)$, by Proposition 3 and a value for $\bar{\mu}_1^n(x)$, the first in line's beliefs after n periods of unsuccessful experimentation given this stationary distribution. As q varies so y_1 and $\bar{\mu}_1$ vary continuously (by (8) and (9)), hence $\bar{\mu}_1^n(x)$ is a continuous function from \mathbb{R}_+ to $[0, 1]$. We will consider the function

$$g(x) := \bar{\mu}_1^{N(x)}(x),$$

that is, the first in line's posterior after $N(x)$ periods of unsuccessful experimentation when the stationary distribution is determined by the strategy $(N(x), q(x))$. This is also

continuous in x and, from, Lemma 4(c), we have $g(x) \rightarrow 0$ as $x \rightarrow \infty$ and $\mu = g(x)$ for any $x < 1$.

If $\mu \leq \bar{\mu}$, it is never optimal for any agent to wait for service to arrive and the equilibrium consists of zero queue lengths. If $\mu > \bar{\mu}$ the continuous function $g(x)$ lies above $f(x)$ for $x = 1$ but $\lim_{x \rightarrow \infty} g(x) - f(x) = 0 - 1 < 0$ by Lemma 4(c). The function $g(x)$ is continuous and the correspondence f has a closed graph on $(0, 1)$, so there must exist $x^* \geq 0$ such that $\bar{\mu}_1^{N(x^*)} = f(x^*)$. If x^* is an integer ($x^* = N(x^*)$) this says, if the first in line had prior $\mu_1^0(x^*)$ and after x^* periods of unsuccessful experimentation it would be optimal for the first in line to leave. If x^* is not an integer ($N(x^*) + 1 > x^* > N(x^*)$) the intersection must occur on a flat portion of f and on such a segment (at the prior $f(x^*)$) the individual is indifferent between engaging in $N(x^*) + 1$ and $N(x^*)$ periods of experimentation. Thus in both cases the first in line's strategy is optimal given the derived beliefs.

Finally, observe that the previous calculations ignored the first in line's priors about timing. By Lemma 2, this moves the function $g(\cdot)$ by at most ϵ and will move the point of intersection, x^* , also by at most ϵ . \square

A.9 Proof of Lemma 6

Assume that $q = 1$. We are interested in the behaviour as $\delta \rightarrow 1$ of equation (6) which determines the equilibrium value of N :

$$N = \mathcal{N}(\bar{\mu}_1^0(N), 1)$$

$$\Leftrightarrow N = \left[(\log(1 - \alpha))^{-1} \log \left(\frac{1 - \bar{\mu}_1^0(N)}{\bar{\mu}_1^0(N)} \frac{\psi(1 - \delta)}{\alpha(\psi^n \delta w - 1)} \right) \right]_+$$

Looking at the continuous version of this (i.e. assuming that both N and M belong to \mathbb{R}_+), and taking the exponential on both sides, we get:

$$(1 - \alpha)^N = \frac{1 - \bar{\mu}_1^0(N)}{\bar{\mu}_1^0(N)} \frac{\psi(1 - \delta)}{\alpha(\psi^n \delta w - 1)}.$$

Simplifying the last term, and substituting the expression for $\bar{\mu}_1^0(N)$, we obtain the following, auxiliary problem:

$$(A.23) \quad (1 - \alpha)^N = \frac{1 - \mu}{\mu} \Delta \frac{x_1(N)}{y_1(N)} =: f(N, \alpha, \delta),$$

where $\Delta := \frac{1 - \delta}{\delta(\alpha w + 1 - \alpha) - 1}$ and $y_1(N)$ and $x_1(N)$ are the stationary probabilities defined in Proposition 3. We can now turn to the proof of Lemma 6.

Proof: (Existence) We begin by establishing that for any $\alpha \in (0, 1)$, or equivalently for any $\phi > 0$, (A.23) admits a solution $N^*(\delta)$ if δ is sufficiently large. We first show that for any δ , as $N \rightarrow +\infty$ the left-hand side of (A.23) tends to zero faster than its right-hand

side. Because $\lim_{N \rightarrow +\infty} N(1 - \alpha)^N = 0$, and $\lim_{N \rightarrow +\infty} \frac{1}{y_1(N)} = \frac{1 - \phi^{M+1}}{1 - \phi}$ for $\phi \neq 1$, and $(M + 1)$ for $\phi = 1$, we obtain that indeed,

$$\lim_{N \rightarrow +\infty} \frac{\frac{1 - \mu}{\mu} \Delta \frac{1}{N y_1(N)}}{(1 - \alpha)^N} = +\infty.$$

On the other hand, for $N = 1$ and for all $\phi > 0$, we have that $y_1(1) = 1$ so that $f(1, \alpha, \delta) = \frac{1 - \mu}{\mu} \Delta$. Since Δ is decreasing in δ for $\delta > (1 - \alpha + \alpha w)^{-1}$, for all $\alpha \in (0, 1)$ and for all $\mu \in (0, 1)$, there exists a $\delta_1 \geq (1 - \alpha + \alpha w)^{-1}$ such that $\forall \delta > \delta_1$, $f(1, \alpha, \delta) < (1 - \alpha)$.

By the continuity of $f(N, \alpha, \delta)$ and $(1 - \alpha)^N$, and by the intermediate value theorem, equation (A.23) therefore admits a solution $N^*(\delta)$ for all $\delta > \delta_1$.

(Lemma 6 (a)) Let $1/2 < \alpha < 1$, or equivalently, let $0 < \phi < 1$. We now turn our attention to the limit of $N^*(\delta)$ as $\delta \rightarrow 1$. For any $\mu \in (0, 1)$,

$$\lim_{\delta \rightarrow 1} f(N, \alpha, \delta) = \frac{1 - \mu}{\mu} \frac{1}{N} \lim_{\delta \rightarrow 1} \frac{\Delta}{y_1(N)}.$$

Unambiguously, we have that $\lim_{\delta \rightarrow 1} \Delta = 0$. Furthermore, for $q = 1$, for all $N \in \mathbb{R}^+$,

$$\lim_{M \rightarrow \infty} y_1(N) = \left[\frac{1}{1 - \phi} \left(1 - (1 + \phi N) \left(\frac{\phi}{1 + \phi} \right)^N \right) \right]^{-1},$$

which belongs to $(0, 1]$ for all $N \in \mathbb{R}^+$ and for all $0 < \phi < 1$. Therefore, $\lim_{\delta \rightarrow 1} f(N, \alpha, \delta) = 0$ for all $N \in \mathbb{R}^+$.

We conclude that, by the continuity in N of both $(1 - \alpha)^N$ and $f(N, \alpha, \delta)$, the solution $N^*(\delta)$ to the auxiliary problem (A.23) tends to $+\infty$ as $\delta \rightarrow 1$, establishing Lemma 6 (b).

(Lemma 6 (b)) Let $0 < \alpha < 1/2$, or equivalently, $\phi > 1$. For any $\mu \in (0, 1)$,

$$\lim_{\delta \rightarrow 1} f(N, \alpha, \delta) = \frac{1 - \mu}{\mu} \frac{1}{N} \lim_{\delta \rightarrow 1} \frac{\Delta}{y_1(N)}.$$

Furthermore, for $q = 1$, we can rewrite:

$$(A.24) \quad y_1(N) = \left[\frac{1}{1 - \phi} [-\phi^{M+1} \kappa + \theta] \right]^{-1},$$

where $\kappa := 1 - \frac{1}{(1 + \phi)^N} \frac{1 - \phi^{N+1}}{1 - \phi}$ and $\theta := 1 - (1 + \phi N) \left(\frac{\phi}{1 + \phi} \right)^N$.

First note that $\lim_{\delta \rightarrow 1} f(1, \alpha, \delta) = 0$ since $y_1(1) = 1$ so that $\lim_{\delta \rightarrow 1} \frac{\Delta}{y_1(1)} = 0$. We now establish that, for all $N > 1$, $\lim_{\delta \rightarrow 1} f(N, \alpha, \delta) = +\infty$. For $N > 1$, we have that $0 < \kappa < 1$ for all $\phi > 0$, and that $|\theta| < \infty$ for all $\alpha > 0$. From (A.24), and since $\lim_{\delta \rightarrow 1} \Delta = 0$ and $\lim_{\delta \rightarrow 1} \phi^{M+1} = +\infty$, we therefore have that:

$$\lim_{\delta \rightarrow 1} \frac{\Delta}{y_1(N)} = -\frac{\kappa}{1 - \phi} \lim_{\delta \rightarrow 1} \Delta \phi^{M+1}.$$

As δ tends to one, we can approximate Δ by $(1 - \delta)[\alpha(w - 1)]^{-1}$. Similarly we approximate ϕ^{M+1} by $\phi^{\frac{1}{1-\delta}}$ so that:

$$\Delta\phi^{M+1} \approx \frac{1 - \delta}{\alpha(w - 1)} \phi^{\frac{1}{1-\delta}},$$

which tends to $+\infty$ when $\delta \rightarrow 1$. This is because $\lim_{\varepsilon \rightarrow 0} \varepsilon \ln \varepsilon = 0$ so that $(\varepsilon \ln \varepsilon + \ln \phi)/\varepsilon$, and therefore $\varepsilon \phi^{1/\varepsilon}$, tend to $+\infty$ as $\varepsilon \rightarrow 0$.

Evaluating the derivative of $f(N, \alpha, \delta)$ at $N = 1$, we obtain:

$$\left. \frac{\partial}{\partial N} f(N, \alpha, \delta) \right|_{N=1} = \phi^{M+1} \left[\frac{(\phi^2 - 1) \ln(1 + \phi) - \phi^2 \ln \phi}{(1 + \phi)(1 - \phi)^2} \right] + \frac{1}{\phi^2 - 1} - \frac{\phi[\ln(1 + \phi) - \ln \phi]}{\phi - 1}.$$

The term multiplying ϕ^{M+1} is positive for $\phi > 1$. Therefore, as $\delta \rightarrow 1$ so that $M \rightarrow +\infty$, we have that $\left. \frac{\partial}{\partial N} f(N, \alpha, \delta) \right|_{N=1} \rightarrow +\infty$.

Finally, observe that for $0 < \alpha < 1$, the left-hand side of (A.23) is strictly between 0 and $1 - \alpha$. From this and the limits of $f(N, \alpha, \delta)$ and $\left. \frac{\partial}{\partial N} f(N, \alpha, \delta) \right|_{N=1}$ as $\delta \rightarrow 1$, we conclude that the solution $N^*(\delta)$ to the auxiliary problem (A.23) is unique and tends to one as $\delta \rightarrow 1$, establishing Lemma 6 (c). Notice that this solution requires that the equilibrium value of q be in $(0, 1)$.

We have just established that there cannot be an equilibrium with $N < M$ for the case $0 < \alpha < 1/2$. This is because $M \rightarrow \infty$ as $\delta \rightarrow 1$, so that for $N > 1$ we have that $y_1(N) = 0$: an individual arriving first in line becomes almost certain that the server is in the bad state. Therefore we cannot have $N > 1$ in equilibrium, since this would give the first in line a payoff of $\delta^N < 1$, and she would be better off balking from the outset.

For $N = 1$ and $q = 1$ the queue can never grow longer than 1 and we have that $y_1(N) = 1$: arriving at the first position in line provides no additional information and each agent optimises on the basis of her prior belief μ . As $\delta \rightarrow \infty$ however, her willingness to experiment grows without bound, and $N = 1$ cannot be an equilibrium either.

For $N = 1$ and $q \in (0, 1)$, μ_1^0 is a continuous, increasing function of q and takes values between 0 and μ . The equilibrium value q^* solves $1 = \mathcal{N}(\mu^*, 1)$ and $U_1(1, \mu^*) = U_1(2, \mu^*)$, where μ^* denotes $\mu_1^0|_{(q, N, M)=(q^*, 1, M)}$.

(Lemma 6 (c)) For $\alpha = 1/2$, we have

$$\frac{x_1(N)}{y_1(N)} = \frac{M(2^{N+1} - 2(N + 1)) + 2^{N+1} - 2 + N(N + 1)}{N2^{N+1}}.$$

so that

$$f(N, 1/2, \delta) = \left(\frac{2^{N+1} - 2 + N(N + 1)}{N2^{N+1}} + \frac{\ln(\delta w)}{\ln(2 - \delta)} \frac{2^{N+1} - 2(N + 1)}{N2^{N+1}} \right) \frac{1 - \mu}{\mu} \frac{2(1 - \delta)}{\delta(w + 1) - 2}$$

We are interested in the limit of $f(N, 1/2, \delta)$ as δ tends to one. Since M tends to infinity, while Δ tends to zero, we simplify:

$$\lim_{\delta \rightarrow 1} f(N, 1/2, \delta) = \lim_{\delta \rightarrow 1} \left\{ \frac{\ln(\delta w)}{\ln(2 - \delta)} \frac{2(1 - \delta)}{\delta(w + 1) - 2} \right\} \frac{1 - \mu}{\mu} \frac{2^{N+1} - 2(N + 1)}{N2^{N+1}}.$$

We obtain the limit of the term in curly brackets using l'Hôpital's rule, so that

$$\lim_{\delta \rightarrow 1} f(1/2, \delta) = \frac{2 \ln w}{w-1} \frac{1-\mu}{\mu} \frac{2^{N+1} - 2(N+1)}{N2^{N+1}}.$$

Therefore, in the limit, the auxiliary problem becomes

$$(1/2)^N = \frac{2 \ln w}{w-1} \frac{1-\mu}{\mu} \frac{2^{N+1} - 2(N+1)}{N2^{N+1}}.$$

A solution N^* to this problem sets the ratio

$$(A.25) \quad \frac{2 \ln w}{w-1} \frac{1-\mu}{\mu} \frac{2^N - (N+1)}{N}$$

equal to one. The ratio's derivative with respect to N ,


$$\frac{2 \ln w}{w-1} \frac{1-\mu}{\mu} \frac{(2^N \ln 2 - 1)N - (2^N - (N+1))}{N^2},$$

is strictly positive whenever $2^N[N \ln 2 - 1] > -1$. This inequality is satisfied for all $N \geq 1$ since the left-hand side is strictly increasing in N , and strictly greater than 2 when $N = 1$. The ratio in (A.25) is therefore strictly increasing in N . Moreover it is easy to see that this ratio is equal to zero when $N = 1$, and tends to infinity when $N \rightarrow +\infty$. We therefore conclude, by the intermediate value theorem, that the solution N^* to the limit of the fixed point problem as $\delta \rightarrow 1$ exists and is unique.

In fact, N^* solves:

$$2^N - 1 = \left(1 + \frac{\frac{1-\mu}{\mu}}{\frac{2 \ln w}{w-1}} \right) N.$$

The right-hand side is increasing in μ . For any $\mu < 1$, the slope of the right-hand side is finite and N^* is finite. For $\mu \rightarrow 1$, $N^* \rightarrow \infty$. For $\mu \rightarrow 0$, N^* solves $2^N - 1 = N$, i.e. $N^* = 1$.



□

A.10 Proof of Lemma 7

Proof: We begin by showing that, based on lemma 6, it is indeed the case that, as $\delta \rightarrow 1$, the equilibrium values of N and M satisfy our assumption that $M > N$. Our assumption is satisfied for $\alpha \leq 1/2$, where we have established that the equilibrium value of N is finite, whereas M goes to infinity. We now show that it is also satisfied when $\alpha > 1/2$.

We begin by showing that $\psi^n \rightarrow 1$ for $n = 1, \dots, N$ as $\delta \rightarrow 1$. For δ close to 1, approximating $y_1(N)$ by:

$$\lim_{M \rightarrow +\infty} y_1(N) = \left[\frac{1}{1-\phi} \left(1 - (1-\phi N) \left(\frac{\phi}{1+\phi} \right)^N \right) \right]^{-1},$$

we can approximate the equilibrium condition, $N = \mathcal{N}(\mu_1^0(N), 1)$, by $(1-\alpha)^N = G(1-\delta)$, for a constant G independent of δ . Because $1-\psi = (1-\delta)\phi\psi$ we have:

$$\log \psi^N = N \log \psi \approx N(\psi - 1) = -N(1-\delta)\phi\psi = -N(1-\alpha)^N \frac{\phi\psi}{G}.$$

(Here we approximate $\log \psi$ by $\psi - 1$ which becomes arbitrarily good as $\delta \rightarrow 1$ or $\psi \rightarrow 1$.) Letting $\delta \rightarrow 1$ or $N \rightarrow \infty$ the right above tends to zero so $\psi^N \rightarrow 1$ as $\delta \rightarrow 1$. Since $1 \geq \psi^n \geq \psi^N$ we have proved our claim. In contrast, $\psi^M \rightarrow 1/w < 1$ as $\delta \rightarrow 1$. A consequence of this is that $M > N$ for all δ large, and our assumption is discharged.

We will now show that deviating from the strategy $\sigma^*(q^*, N^*, M^*)$ is suboptimal for the n^{th} in line, where $n = 2, \dots, M+1$. First observe that reneging when the first in line reneges is always optimal for arrival $n > 1$. When the first in line reneges the n^{th} in line's posterior on the queue state equals the first in line's (as their information sets are nested). If it were optimal for the first in line to exit and get the payoff 1 those behind her in the queue strictly prefer to exit as their value to waiting in the good state is strictly less than the first in line's value. Second, observe that those who arrive in place $M^* \geq n > N^* + 1$ strictly prefer to wait for service because they know the server is good. Third, observe that $M^* = \lfloor \mathcal{M} \rfloor$ ensures that it is suboptimal to join longer queues.

It remains to check that no n^{th} in line (for $n = 2, \dots, N^* + 1$) prefers to renege before they observe the first in line reneging. As before, we begin by assuming the posteriors are determined ignoring the prior information on the calendar date. We will show that each possible deviation from the strategy $\sigma^*(q^*, N^*, M^*)$ reduces the expected payoff of an individual arriving n^{th} in line.

First, we describe the n^{th} -in-line's payoff, for $n = 2, \dots, N^* + 1$, from deviating from her equilibrium strategy and reneging before she observes the first in line renege. The n^{th} in line may have to wait until just before the $N^* + 2^{\text{st}}$ arrival to observe the first in line renege, so early renege can occur after $m = 0, 1, \dots, N^* - n + 1$ periods, although the first in line may also renege in the last of these periods with positive probability. In (4) we defined the n^{th} -in-line's expected value of reneging after m unsuccessful service events (and waiting if service is observed) assuming there was no social learning.

$$U_n(m, \bar{\mu}_n^0) = \bar{\mu}_n^0 \left[\frac{1 - \bar{\mu}_n^0}{\bar{\mu}_n^0} \delta^m + \psi^n w \delta - (1-\alpha)^m \delta^m (\psi^n w \delta - 1) \right]$$

If the n^{th} in line reneges in periods $m = 0, 1, \dots, N^* - n$ there is no social learning, because the first in line never reneges until just before the line reaches length $N^* + 1$. Thus, $U_n(m, \bar{\mu}_n^0)$ equals the n^{th} -in-line's expected payoff if she chooses to deviate from her equilibrium strategy and renege at $m = 0, 1, \dots, N^* - n$. If the n^{th} in line chooses to renege at $m = N^* - n + 1$ and the first in line has not, then there is social learning because the first in line's decision whether to renege or not is informative about her posterior belief. The social learning (from seeing the first in line not renege) would lead the n^{th} in line to revise upwards the probability she attaches to the server being good. Thus the expression above overestimates the payoff from renegeing at $m = N^* - n + 1$ by omitting the benefit of this social learning. Hence $U_n(m, \bar{\mu}_n^0)$, for $n = 2, \dots, N^* + 1$ and $m = 0, \dots, N^* - n + 1$, is an overestimate of the payoff from early renegeing at this equilibrium.

To verify equilibrium we will show that $U_n(m, \bar{\mu}_n^0)$, for $n = 2, \dots, N^* + 1$ and $m = 0, \dots, N^* - n + 1$, is less than what the n^{th} in line expects to get by renegeing if and only if the first in line does. We will use a lower bound for the payoff to abiding by the equilibrium by assuming that the first in line experiments for $N^* + 1$ periods with certainty, and then reneges (rather than renegeing with positive probability after N^* unsuccessful service events); that is, we consider the case where $q = 0$.

Assume that the server is in the good state and let A_n be the n^{th} in line's expectation of the payoff she will obtain once the queue has reached length $N^* + 1$, and the first in line's behaviour (renege or not) reveals her posterior belief ($\mu_1^{N^*+1}$ or 1 respectively). By Proposition 3 (and ignoring the prior on timing), the individual who joined the queue at the n^{th} position attaches probability $y_1(1 - \alpha)^{n-1}/y_n$ to the to the first-in line never having observed service. This is, therefore, the probability that the n^{th} in line attaches to her renegeing together with the first in line once the queue reaches length $N^* + 1$, in which case the n^{th} in line's payoff is 1. She attaches the complementary probability to the first in line not renegeing once the queue reaches length $N^* + 1$. In that case, the n^{th} in line's payoff is $\psi^n \delta w$, by (1).

Thus

$$A_n = \psi^n \delta w \left(1 - \frac{(1 - \alpha)^{n-1} y_1}{y_n} \right) + \frac{(1 - \alpha)^{n-1} y_1}{y_n} \geq 1.$$

If the n^{th} in line abides by the equilibrium and waits for $N^* - n + 2$ unsuccessful service events so as to herd on the first in line's behaviour when the queue is length $N^* + 1$, then her payoff is:

$$U_n^* = \bar{\mu}_n^0 \left[\frac{1 - \bar{\mu}_n^0}{\bar{\mu}_n^0} \delta^{N^* - n + 2} + \psi^n w \delta - (1 - \alpha)^{N^* - n + 2} \delta^{N^* - n + 2} (\psi^n w \delta - A_n) \right].$$

A sufficient condition for early renegeing to be suboptimal for the n^{th} in line is $U_n^* \geq U_n(m, \mu_n^0)$ for all $n = 2, \dots, N^* + 1$ and $m = 0, \dots, N^* - n + 1$. Substituting from above, early renegeing is suboptimal if

$$(A.26) \quad \frac{1 - \bar{\mu}_n^0}{\bar{\mu}_n^0} (\delta^m - \delta^{N^* - n + 2}) \leq (1 - \alpha)^m \delta^m (\psi^n w \delta - 1) - (1 - \alpha)^{N^* - n + 2} \delta^{N^* - n + 2} (\psi^n w \delta - A_n)$$

for all $n = 2, \dots, N^* + 1$ and $m = 0, \dots, N^* - n + 1$. Substituting the value of A_n we

obtain

$$\frac{1 - \bar{\mu}_n^0 \delta^m - \delta^{N^*-n+2}}{\bar{\mu}_n^0 \psi^n w \delta - 1} \leq (1 - \alpha)^m \delta^m - \delta^{N^*-n+2} \frac{(1 - \alpha)^{N^*+1} y_1}{y_n}.$$

Now substituting from (14), the definition of $\bar{\mu}_n^0$, we get the equilibrium condition

$$\frac{1 - \mu \delta^m - \delta^{N^*-n+2}}{N^* \mu \psi^n w \delta - 1} \leq y_n \delta^m (1 - \alpha)^m - y_1 (1 - \alpha)^{N^*+1} \delta^{N^*-n+2}.$$

Finally, dividing through by $\delta^m (1 - \alpha)^{m+n}$, writing $s = m + n$ and substituting the expressions from Proposition 3 for y_1 and y_n gives the sufficient condition for equilibrium:

$$(A.27) \quad \frac{1 - \mu}{BN^* \mu} \frac{1 - \delta^{N^*-s+2}}{(1 - \alpha)^s} \leq (\psi^n w \delta - 1) \left(\frac{\phi^{n-1} - k_{N^*}}{(1 - \alpha)^n} - (1 - k_{N^*})(1 - \alpha)^{N^*+1-s} \delta^{N^*-s+2} \right),$$

for $s = 2, \dots, N^* + 1$ and $n = 2, \dots, s$.

We now show that for δ sufficiently close to one (A.27) holds. First observe that (A.27) also holds for $n = 1$ and $s = 1, \dots, N^* + 1$. since, by the construction of the equilibrium N^* , it is optimal for the first in line to experiment for $N^* + 1$ periods and then renege. There is no social learning for the first in line and the indifference that occurs if $q^* > 0$ ensures that there are no estimates, so in this case the expressions are exact.

Now we consider (A.27) for $n > 1$. The term $(\phi^{n-1} - k_{N^*})/(1 - \alpha)^n$ in (A.27) increases in n for $\alpha > 1/2$ or $k_{N^*} < 0$ ($\alpha < \alpha_{N^*}^*$).¹⁷ This is because

$$\frac{\phi^n - k_{N^*}}{(1 - \alpha)^{n+1}} - \frac{\phi^{n-1} - k_{N^*}}{(1 - \alpha)^n} = \frac{\alpha}{(1 - \alpha)^{n+1}} (\phi^{n+1} - k_{N^*}^*).$$

When $k_{N^*}^* < 0$ and $\alpha_{N^*}^* > \alpha$ the right-hand-side above is greater than unity. When $\alpha > 1/2$ the right-hand-side above is increasing in n so is minimised at $n = 2$. Thus we get

$$\frac{\phi^n - k_{N^*}}{(1 - \alpha)^{n+1}} - \frac{\phi^{n-1} - k_{N^*}}{(1 - \alpha)^n} \geq \begin{cases} 1, & \alpha < \alpha_{N^*}^*; \\ (\phi^3 - k_{N^*}) \phi^{-3} \alpha^{-2}, & \alpha > 1/2. \end{cases}$$

Hence the first term in n is approximately constant for δ large, by our initial argument, while the second term strictly increases in n by an amount independent of δ . In combination, therefore, as n increases the RHS of (A.27) increases in n , for all s . As the inequality holds for $n = 1$ it holds for all $n > 1$ too. Hence we have shown that no later arrival can benefit from early entry.

Finally, note that this finite set of inequalities hold strictly for the beliefs $\bar{\mu}_n^0$. By Lemma (2), therefore, they hold for the beliefs μ_n^0 for ν sufficiently small. As ν depends on the number of states in the Markov process and N^* increases in δ how small ν must be depends on δ . \square

¹⁷This was defined in the proof of Proposition 3.

B Appendix

Here we describe the solution to the team problem. The first step is to describe the stationary distribution of queue lengths in the good state when the team adopts the strategy described in Section 4. This is a simplified version of the calculation in the first Appendix and yields the distribution:

$$(B.1) \quad z_n = \frac{1 - \phi}{1 - \phi^{M^\dagger+1}} \begin{cases} \phi^{n-1}, & n = 1, 2, \dots, M^\dagger - 1; \\ \phi^{M^\dagger-1}/\alpha, & n = M^\dagger. \end{cases}$$

The next step is to describe the team's time-average utility. By the ergodic theorem, z_n is also the time-average of the number of agents arriving at the queue at the n^{th} position for $n < M^\dagger$. Such an agent will get utility $\psi^n \delta w$ if she waits for service. The average number of agents arriving at the queue at the M^{th} position can be found by considering their history. They must have joined a line of length $M^\dagger - 1$, which occurs with stationary probability $(1 - \alpha)z_{M^\dagger-1} + (1 - \alpha)\alpha z_{M^\dagger}$. (This is because an agent can find $M^\dagger - 1$ agents in the line if no service occurred or if one unit of service occurred and the line was previously length M^\dagger). These agents get utility $\psi^{M^\dagger} \delta w$. Finally, an agent will balk and get utility one if she arrives at the queue at the $M^\dagger + 1^{\text{st}}$ position. This occurs with probability $(1 - \alpha)z_{M^\dagger}$. Hence the time average welfare of the team is

$$W(M^\dagger) := \sum_{n=1}^{M^\dagger-1} z_n \psi^n \delta w + \psi^{M^\dagger} \delta w (1 - \alpha)(z_{M^\dagger-1} + \alpha z_{M^\dagger}) + (1 - \alpha)z_{M^\dagger}.$$

Substituting for the stationary distribution and evaluating the summations gives

$$W(M^\dagger) = \frac{1 - \phi}{1 - \phi^{M^\dagger+1}} \left[\psi \delta w \left(\frac{1}{1 - \phi \psi} - \phi^{M^\dagger} \psi^{M^\dagger} \frac{2 - \phi \psi}{1 - \phi \psi} \right) + \phi^{M^\dagger} \right].$$

A tedious calculation can then be performed to show:

$$W(M + 1) - W(M) = \phi^M \frac{1 - \phi}{1 - \phi^{M+1}} \left[(2 - \phi \psi) \psi^{M+1} \delta w - (1 - \phi) - \phi W(M + 1) \right].$$

The socially optimal value of M^\dagger is the smallest value of M for which the above expression is negative.