

Supplement to “Strategic experimentation in queues”

(*Theoretical Economics*, Vol. 14, No. 2, May 2019, 647–708)

MARTIN W. CRIPPS

Department of Economics, University College London

CAROLINE D. THOMAS

Department of Economics, University of Texas at Austin

APPENDIX B

B.1 *Definition of a strategy*

We begin by defining the set of possible histories of play an agent can observe conditional on her not having been served and not having exited: these are the only circumstances under which she has to act. Recall that \mathbb{N}^0 is the set of nonnegative integers.

On arrival at the queue, the agent observes the position at which she arrives at the queue, which is an element of

$$H_0 := \mathbb{N}^0.$$

The agent then witnesses her first service stage, where she observes whether service occurs. The number of agents remaining in the queue after that service stage is sufficient information to determine whether service occurred. (The difference determines the number who were served.) Thus the set of histories after the first service stage is

$$H_{10} := \mathbb{N}^0 \times H_0.$$

The agent then witnesses her first exit stage, where she observes which agents in the queue choose to renege (E) on the queue and which continue (C) queuing. Recall that at every exit *stage* we allow multiple *rounds* of exit. If the queue has length n at the beginning of the first exit stage, the agent’s observation after the first round of exit is an element of $\{E, C\}^n$. Thus the set of histories after the first round of the first exit stage is

$$H_{11} := \bigcup_{n=0}^{\infty} \{E, C\}^n \times H_{10}.$$

By Lemma 1, \mathcal{M} provides an upper bound on the possible queue lengths when agents are rational. We therefore allow for a total of \mathcal{M} rounds of exit. This ensures that for any rational queue length, there are sufficient opportunities for all agents in the

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

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

queue to exit at this stage, even if only one agent were to exit at each round.¹ Thus the set of histories after the m th round of the first exit stage is

$$H_{1m} := \left(\bigcup_{n=0}^{\infty} \{E, C\}^n \right)^m \times H_{10}, \quad \text{for } m = 2, \dots, \mathcal{M}.$$

Finally, the agent witnesses her first arrival stage, at which one additional agent arrives and decides whether to join the queue or to balk. This determines a new number of agents in the queue. Thus the set of histories after the first arrival stage is

$$H_1 := \mathbb{N}^0 \times H_{1\mathcal{M}}.$$

This is the set of all possible histories an agent could observe after being present in the queue for one complete period of our game.

We now recursively define H_t , the set of histories an agent could have observed after being present in the queue for t complete periods:

$$H_t := \mathbb{N}^0 \times \left(\bigcup_{n=0}^{\infty} \{E, C\}^n \right)^{\mathcal{M}} \times \mathbb{N}^0 \times H_{t-1}.$$

We define the sets H_{tm} for $m = 0, \dots, \mathcal{M}$ and $t > 1$ analogously with the sets H_{1m} defined above.

A behavior strategy for an agent gives a probability distribution over her choices to either exit (E) or continue in the queue (C) when she arrives at the queue, as well as at the \mathcal{M} exit rounds in each period after she joins the queue. Observe that before choosing her action in the m th exit round of her t th exit stage, the agent observes a history in $H_{t(m-1)}$. Thus, the set of histories at which the agent is called upon to act is

$$H := H_0 \cup \left(\bigcup_{t=0}^{\infty} \bigcup_{m=0}^{\mathcal{M}-1} H_{tm} \right),$$

and a behavior strategy for the agent is a mapping

$$\sigma : H \rightarrow \Delta(\{E, C\}).$$

¹Observe that our restriction to \mathcal{M} rounds of exit does not mean that we allow only for rational queue lengths to arise. However, it does mean that should queue lengths longer than \mathcal{M} arise, then the continuation play may be constrained by the \mathcal{M} -rounds exit-stage protocol described here. To illustrate, suppose the queue has length $\mathcal{M} + k$ and only one agent exits at each exit round. Then k agents will not have exited at the end of the current exit stage (even though they might have exited if there had been $\mathcal{M} + k$ rounds), and will, therefore, have to wait for the next period's exit stage. Allowing \mathcal{M} rounds of exit ensures that the exit-stage protocol imposes no restrictions on play whenever the queue length is rational. In particular, observe that two rounds of exit would have been sufficient in our equilibria. At the first round, the first in line chooses C or E . At the second round, all other agents in the queue herd on the first in line's chosen action.

B.1.1 Definition of the equilibrium strategy To formally define the strategy $\sigma^*(q, N, M)$ given in Definition 1, we introduce the following notation. Let n_t denote the queue length at history $h_t \in H_t$ and let n_{tm} denote the queue length at history $h_{tm} \in H_{tm}$. Let $a_{tm} \in \{E, C\}^{n_{t(m-1)}}$ denote the action profile played at the m th exit round of stage t and let $a_{tm}(i)$ denote its i th coordinate. Observe that $a_{tm}(1) = E$ indicates that at the m th exit round of stage t , the first agent in line reneges. The strategy $\sigma^*(q, N, M)$ is then defined as follows (for $q = 1$).

First, upon arriving at the queue, the agent joins the queue if and only if she is at most M th in line. Thus,

$$\sigma(h_0) = \begin{cases} E & \text{if } n_0 > M, \\ C & \text{if } n_0 \leq M. \end{cases}$$

For $m = 0, \dots, M - 1$ and $t \geq 0$, we now describe the agent's strategy at a history $h_{tm} \in H_{tm}$. If there exists a date $0 < \tau \leq t$ such that $n_{\tau 0} \neq n_{(\tau-1)}$, then the agent chooses C . Thus, once the agent observes service, she never reneges.

Conversely, if $n_{\tau 0} = n_{(\tau-1)}$ for every $0 < \tau \leq t$ so that the agent has never observed service, her behavior depends on the position at which she joined the queue and on the behavior of other agents in the queue.

If $n_0 = 1$ so that the agent joined the queue at the first position, she reneges after having observed the N th service failures, i.e., when $t = N$ and $m = 0$. She continues for every $t < N$ (and $m = 0, \dots, M - 1$).

If $n_0 > 1$ so that there were agents ahead of her in the queue when she joined, then she reneges on the queue if and only if the first in line does and in the same period as the first in line. We distinguish two cases.

- If she observes the first in line renege at the m th round of the t th exit stage, for $m = 1, \dots, M - 1$, then the agent exits at the $(m + 1)$ th round of the t th exit stage. That is, the agent plays E if h_{tm} is such that $m = 1, \dots, M - 1$ and $a_{tm}(1) = E$.
- If she observes the first in line renege at the last (M th) round of the $(t - 1)$ th exit stage, then the agent exits at the first round of the t th exit stage, provided there was no service at the t th service stage. That is, the agent plays E if h_{tm} is such that $m = 0$, $a_{(t-1)M}(1) = E$, and $n_{(t-1)} = n_{t0}$ (no service at the t th service stage).

We can therefore partition the histories $H \setminus H_0$ into two subsets: H^E , the set of histories at which the agent reneges on the queue, and its complement, $H \setminus (H_0 \cup H^E)$, the set of histories at which the agent continues queuing. We then have

$$\sigma(h_{tm}) = \begin{cases} E & \text{if } h_{tm} \in H^E, \\ C & \text{otherwise,} \end{cases}$$

where $H^E \subseteq H$ are the histories h_{tm} such that $n_{\tau 0} = n_{(\tau-1)}$ for every $0 < \tau \leq t$ and such that any of the following three conditions is satisfied:

- We have $n_0 = 1$, $t = N$, and $m = 0$.

- We have $n_0 > 1$, $m = 1, \dots, \mathcal{M} - 1$, and $a_{tm}(1) = E$.
- We have $n_0 > 1$, $m = 0$, $a_{(t-1),\mathcal{M}}(1) = E$, and $n_{(t-1)} = n_{t0}$.

B.2 The good state stationary measure on the interval $[1, \phi_N^*]$

When $1 < \phi < \phi_N^*$. For these values of ϕ , 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 $N < M$. In these cases, 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 for $N < M$, the stationary measure is U-shaped. It jumps down at $N + 1$ and $N + 2$, and then increases over the range $N + 2 \leq n \leq M$, as illustrated in Figure B1.




FIGURE B1. The stationary measure of the queue length conditional on the server being in the good state with $1 < \phi < \phi_N^*$ under perfect revelation (left panel) and imperfect revelation (right panel).

Knife-edge cases: $\phi = 1$ and $\phi = \phi_N^$.* For $\phi = 1$, service is exactly as fast as arrivals. If agents never renege, but waited in line until served, then every queue length would be equally likely. For $N < M$, this is in fact the case when $n \geq N + 2$, and the stationary measure is uniform over these values. When $1 \leq n \leq N$, the renewal effect is a force for emptying the queue and the stationary measure is linearly decreasing over these values. The downward steps at $n = N + 1$ and $n = N + 2$ remain. For $M \leq N$, the renewal effect is always at play, and the stationary measure is linearly decreasing in n for $1 \leq n \leq M$ and has a downward step at $n = M + 1$.

For $\phi = \phi_N^*$, service is exactly slow enough to offset the renewal effect. Consequently, for $M \leq N$, the stationary measure is uniform for $1 \leq n \leq N$. Without the renewal effect, the queue tends to fill up, so that the stationary measure is increasing for $n \geq N + 2$. The downward steps at $n = N + 1$ and $n = N + 2$ remain. For $M \leq N$, the stationary measure is uniform for $1 \leq n \leq M$ and has a downward step at $n = M + 1$. These cases are illustrated in Figure B2.

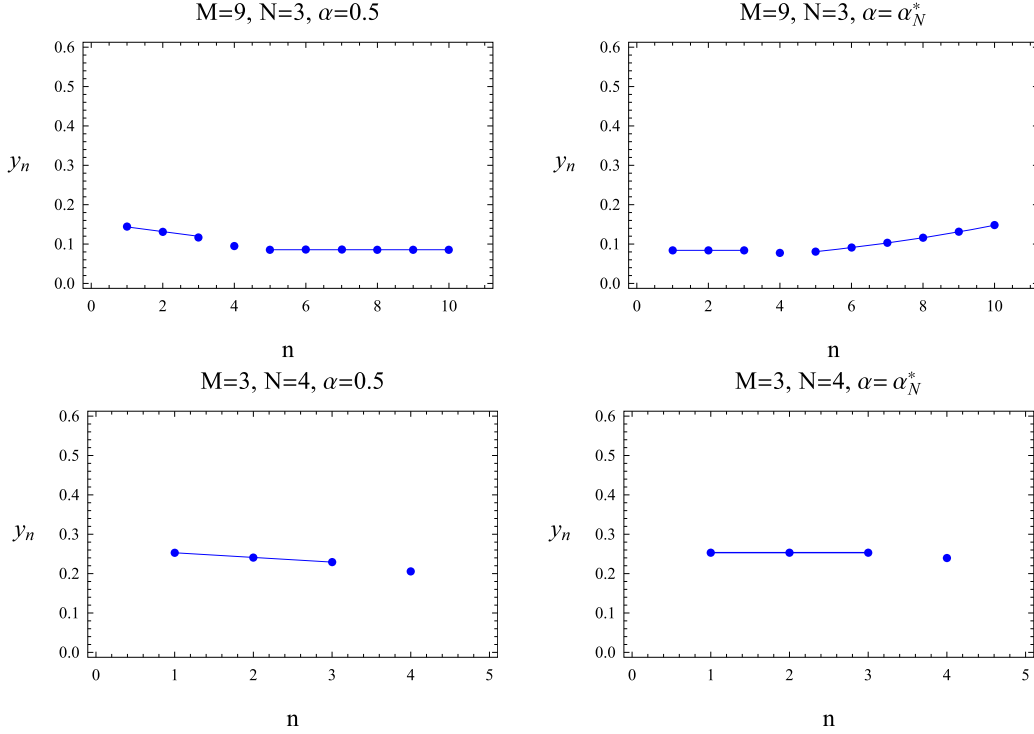


FIGURE B2. The stationary measure of the queue length conditional on the server being in the good state when ϕ takes the values 1 and ϕ_N^* under perfect revelation (upper two panels) and imperfect revelation (lower two panels).

B.3 Stationary measure for $\alpha = 1/2$

LEMMA B.1. *Let $\alpha = 1/2$. For $N < M$, the stationary measure of queue lengths in the good state 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 measure of queue lengths in the good state 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.10) 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.10), on this equation, we solve for H and obtain

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

Substituting into the expressions for y_N , y_{N+1} , and y_{N+2} in (A.10), 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.10), gives $y_M = y_{M+1}$, and from the expression for y_n when $N + 2 < n < M + 1$ in (A.10), 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 B.1. (The case $N \geq M$ can be analyzed in a similar fashion.) \square

B.4 Equilibria with multiple herding leaders

We begin by giving a formal definition of a *herding strategy* $\check{\sigma}$, that is, a pure strategy with multiple herding leaders. In this strategy, the individuals at positions $1, \ell_2, \dots, \ell_C$ in the queue are the *herding leaders*. Only herding leaders may autonomously leave the queue. As under σ^* , under $\check{\sigma}$ the first in line autonomously reneges after $L_1 = N$ periods without service. The herding leader at position ℓ_c in the line autonomously reneges after observing $L_c \in \mathbb{N}$ periods without service or reneges if someone ahead of her does. All remaining agents are *herding followers*. They renege if and only if they observe someone ahead of them renege, but stay in line if they ever observe service.

DEFINITION B.1. The strategy with $C \geq 2$ herding leaders is $\check{\sigma}(N, (\ell_c)_{c=2}^C, (L_c)_{c=2}^C, M)$ with $1 < \ell_c < \ell_{c+1}$ for all $c \geq 2$, $L_2 < N - \ell_2 + 1$, and $L_c < L_{c-1} - \ell_c + 1$ for all $c \geq 3$:

- Upon arriving at the queue, an individual joins the queue if and only if she is at most M th in line.
- Once in the queue, if she observes service, she never reneges.
- Conditional on not observing service, the following scenarios occur:
 - If she joined the queue at the first position then she is the first herding leader. She does not renege for the first $N - 1$ periods and reneges at the exit stage of the N th period.
 - For each $c = 2, \dots, C$, if she joined the queue at the ℓ_c th position, then she is the c th herding leader. If at any point she observes an agent ahead of her renege, she reneges in the same period as that agent. Otherwise, she does not renege for the first $L_c - 1$ periods and reneges at the exit stage of the L_c th period.
 - If she joined the queue at any of the remaining positions, then she is a herding follower. She reneges if and only if an agent ahead of her does and in the same period as that agent.

As one might expect, the stationary distributions under these strategies are quite complex. Figure B3 illustrates the behavior of the queue at a bad server when all agents follow the strategy $\check{\sigma}(N, (\ell_c)_{c=2}^C, (L_c)_{c=2}^C, M)$. In the first panel, $C = 2$ and there are two herding leaders: the first in line, who autonomously reneges on the queue if the 21 first service opportunities she observes are unsuccessful, and the third in line, who autonomously reneges on the queue if the 8 first service opportunities she observes are unsuccessful. Thus, the strategy in the first panel has $(N, \ell_2, L_2) = (21, 3, 8)$ (and $M \geq 10$). In the second panel, $C = 3$ and there is an additional herding leader: the fifth in line, who autonomously reneges on the queue if the first 4 service opportunities she observes are unsuccessful. Thus, the strategy in the second panel has $(N, (\ell_2, \ell_3), (L_2, L_3)) = (21, (3, 5), (8, 4))$ (and $M \geq 8$). Observe that in both examples, there are three individuals joining the queue at the third position in line behind a given first in line, and that in the second example, there are five individuals joining the queue at the fifth position behind a given first in line.




FIGURE B3. The behavior of the queue at a bad server under the strategy $\check{\sigma}(21, 3, 8, M)$ for $M > 10$ (first panel) and $\check{\sigma}(21, (3, 5), (8, 4), M)$ for $M > 8$ (second panel). (These strategies need not be equilibrium strategies.)

Under the strategy $\check{\sigma}$, we assume that at a bad server there are always at least two² instances of the ℓ_c th in line behind a given first in line, for every $c = 2, \dots, C$. At a good server, there is no upper bound on the possible number of instances of the ℓ_c th in line behind a given first in line. Crucially, an individual joining the queue at the ℓ_c th position does not know whether she is the first, second, \dots instance of the ℓ_c th in line behind a given first in line. She learns both about this and the server state while waiting in line.

The next lemma provides conditions under which a strategy with more than one herding leader cannot be an equilibrium. The proof is given in [Appendix B.5](#).

LEMMA B.2. *Consider the strategy $\check{\sigma}(N, (\ell_c)_{c=2}^C, (L_c)_{c=2}^C, M)$ with $C \geq 2$. If $\alpha \geq 1/2$, there exists a $\check{\delta} < 1$ such that for every $\delta > \check{\delta}$, this strategy does not constitute an equilibrium.*

²If that is not the case, then the herding learner in question reneges on the queue when another agent ahead of her does. She is then effectively a herding follower.

Lemma B.2, together with Lemma 6, implies that for $\alpha \geq 1/2$, there exist values of δ sufficiently close to 1 at which any given strategy with more than one herding leader cannot be an equilibrium, but at which there exists an equilibrium with perfect revelation at which the first in line is the unique herding leader.

The intuition for **Lemma B.2** is as follows. Let \check{N} denote the optimal duration of experimentation for the first in line and let \check{L}_c denote the optimal duration of experimentation for the ℓ_c th in line when agents' beliefs are consistent with the stationary distributions generated by the strategy $\check{\sigma}(\check{N}, (\ell_c)_{c=2}^C, (\check{L}_c)_{c=2}^C, M)$. Consider the second herding leader, the ℓ_2 th in line, and the individual just ahead of her, the $\ell_2 - 1$ th in line. In equilibrium, the $\ell_2 - 1$ th in line finds it optimal to herd on the first in line. Given the posterior belief she forms upon arriving at the queue at the $\ell_2 - 1$ th position (based on the stationary probabilities of arriving at that position in each state), she must find it optimal to remain in line for at least $\check{N} - \ell_2$ periods, giving the first in line enough time to complete her \check{N} periods of experimentation. In contrast, the ℓ_2 th in line must find it optimal to renege autonomously after only $\check{L}_2 < \check{N} - \ell_2 - 1$ periods, conditional on not observing service. The proof of the lemma argues that these two requirements contradict each other when $\alpha \geq 1/2$ and δ is sufficiently close to 1. The reason is that the individual arriving at position ℓ_2 is too optimistic to leave the queue so much earlier than the individual at position $\ell_2 - 1$. Our proof does not require calculating the stationary distribution explicitly. Instead we use bounds on the stationary distribution that arise under the strategy $\check{\sigma}$.

This follows from the peculiar aspect social learning takes for the ℓ_2 th in line. If she were the k th instance of the ℓ_2 th in line behind a given first in line who is as yet uninformed, then after $\check{N} - (k - 1)\check{L}_2 - \ell_2 + 1$ service failures she should observe the uninformed first in line renege. Thus, if the first in line does not renege at that point, the ℓ_2 th in line learns that she is not the k th instance of the ℓ_2 th in line behind an uninformed first in line. As a result, her posterior belief that the server is good does not fall commensurately in that period with that of the $\ell_2 - 1$ th in line. (For some parameter values her posterior may even jump up.) This increase in the ℓ_2 th in line's optimism relative to that of the $\ell_2 - 1$ th in line's implies that the ℓ_2 th in line cannot find it optimal to renege after observing only \check{L}_2 service failures if the $\ell_2 - 1$ th in line accepts to observe $\check{N} - \ell_2$ service failures without renegeing.

B.5 Proof of Lemma B.2

PROOF. Assume, by way of contradiction, that the strategy $\check{\sigma}(\check{N}, (\ell_c)_{c=2}^C, (\check{L}_c)_{c=2}^C, M)$ with $C \geq 2$ herding leaders constitutes a symmetric equilibrium. We concentrate on the first two herding leaders: the first in line and the ℓ_2 th in line. To lighten notation, in this appendix we use ℓ and \check{L} for ℓ_2 and \check{L}_2 , respectively.

In addition we define the following notation. Let $j = 1, \dots, \check{L}$ and $k = 2, 3, \dots$ satisfy

$$\check{N} = (k - 1)\check{L} + \ell + j - 1. \quad (\text{B.1})$$

The variables k and j admit the following interpretation. At a bad server or at a good server at which the first in line is uninformed, there are k instances of the ℓ th in line for

every instance of the first in line. Furthermore, the k th instance of the ℓ th in line will learn the first in line's information after observing j unsuccessful service events, as this coincides with the first in line completing her \check{N} periods of experimentation.

Finally, we let \check{y}_n denote the stationary probability of arriving at the queue at the n th position in line when the server is good under $\check{\sigma}$, our candidate equilibrium strategy with multiple herding leaders.

We show that for certain parameter values, if it is optimal for the first in line to experiment for \check{N} periods and it is optimal for the $\ell - 1$ th to herd on the first in line, then it cannot be optimal for the ℓ th in line to experiment for \check{L} periods. Our argument is invariant to the presence of further herding leaders.

Begin by considering the first in line. Under $\check{\sigma}$, the stationary probability of arriving at the first position at a bad server is $1/\check{N}$; at a good server, it is \check{y}_1 . In equilibrium, the first in line finds it optimal to experiment for \check{N} periods. Furthermore, she does not learn anything from observing the behavior of others in the queue. Therefore, using the individual threshold defined in (10), \check{N} is determined by the relationship

$$\frac{\mu\check{N}\check{y}_1(1-\alpha)^{\check{N}-1}}{1-\mu} \geq \frac{\underline{\mu}_1}{1-\underline{\mu}_1} > \frac{\mu\check{N}\check{y}_1(1-\alpha)^{\check{N}}}{1-\mu}. \quad (\text{B.2})$$

Now consider an individual arriving at the n th position in line, for $n = 2, \dots, \ell - 1$. The probability of arriving at that position is \check{y}_n at a good server and $1/\check{N}$ at a bad server. The n th in line cannot learn anything from the behavior of others except the first in line. Thus, as long as the first in line does not renege and conditional on no service, the likelihood ratio of the n th in line's posterior belief follows the path

$$\frac{\mu\check{N}\check{y}_n}{1-\mu} > \frac{\mu\check{N}\check{y}_n(1-\alpha)}{1-\mu} > \dots > \frac{\mu\check{N}\check{y}_n(1-\alpha)^{\check{N}-n+1}}{1-\mu}.$$

Equilibrium requires that the n th in line does not want to renege before the first in line has completed her \check{N} periods of experimentation. Equivalently,³ for all $n = 1, \dots, \ell - 1$,

$$\frac{\mu\check{N}\check{y}_n(1-\alpha)^{\check{N}-n}}{1-\mu} \geq \frac{\underline{\mu}_n}{1-\underline{\mu}_n}. \quad (\text{B.3})$$

Let us now consider an individual's inference upon arriving at the queue at the ℓ th position. That individual does not know whether or how many other individuals have already arrived at the ℓ th position behind the current first in line. In particular, conditional on the server being bad, the individual believes she could equiprobably be the first, second, \dots , k th instance of the ℓ th in line behind that given first in line. Thus, the probability she attaches to arriving at the ℓ th position at a bad server is k/\check{N} .

³If $\frac{\mu\check{N}\check{y}_n(1-\alpha)^{\check{N}-n+1}}{1-\mu} < \frac{\underline{\mu}_n}{1-\underline{\mu}_n}$, i.e., if the n th in line would like to renege at the exit stage following the $(N^* - n + 1)$ th failure based on her private learning alone, we assume that she still waits to observe the first in line's behavior.

A similar reasoning applies at a good server if the first in line is uninformed. Consequently, the stationary probability of arriving at the queue at the ℓ th position in line is

$$\check{y}_\ell = b_0 + b_1 + \dots + b_k,$$

where b_0 is the stationary probability of arriving at the ℓ th position at a good server when the first in line has already observed service, and b_m for $m = 1, \dots, k$ is the stationary probability of arriving as the m th instance of the ℓ th in line at a good server when the first in line is uninformed. Equivalently, b_m is the probability that the current first in line joined the queue at the first position and has since observed $\ell - 1 + (m - 1)\check{L}$ unsuccessful service events: $b_m = \check{y}_1 (1 - \alpha)^{\ell - 1 + (m - 1)\check{L}}$ for $m = 1, \dots, k$.

Now suppose that an individual knew that she joined the queue as the k th instance of the ℓ th in line (this is not the case in equilibrium). She would only need to wait j periods to obtain the first in line's information. If at that point the first in line reneges, all individuals in the queue learn that she has observed \check{N} consecutive service failures. If she does not renege, an individual who knew she was the k th instance of the ℓ th in line would learn that the first in line is informed and, hence, that the server is good.

Likewise, the first in line's behavior is also informative for an individual who joins the queue at the ℓ th position, but *does not know* which instance of the ℓ th in line she is. If she observes j failures and the first in line *not* renegeing, the ℓ th in line learns that (i) she is not the k th instance of the ℓ th in line behind an uninformed first in line at a good server (probability b_k) and (ii) she is not the k th instance of the ℓ th in line behind the first in line at a bad server (probability $1/\check{N}$).

Thus, for the first $j - 1$ failures, the likelihood ratio of the ℓ th in line's posterior belief follows the path

$$\frac{\mu \check{N} \check{y}_\ell}{(1 - \mu)k} > \frac{\mu \check{N} \check{y}_\ell (1 - \alpha)}{(1 - \mu)k} > \dots > \frac{\mu \check{N} \check{y}_\ell (1 - \alpha)^{j-1}}{(1 - \mu)k},$$

and at the j th failure, if the first in line does not renege, the ℓ th in line updates the likelihood ratio of her posterior belief to

$$\frac{\mu \check{N} (\check{y}_\ell - b_k) (1 - \alpha)^j}{(1 - \mu)(k - 1)}.$$

We conclude that in equilibrium, since she must find it optimal to autonomously exit the queue after \check{L} periods without service, the likelihood ratio of the ℓ th in line's posterior belief must satisfy

$$\frac{\mu \check{N} (\check{y}_\ell - b_k) (1 - \alpha)^{\check{L}-1}}{(1 - \mu)(k - 1)} \geq \frac{\underline{\mu}_\ell}{1 - \underline{\mu}_\ell} > \frac{\mu \check{N} (\check{y}_\ell - b_k) (1 - \alpha)^{\check{L}}}{(1 - \mu)(k - 1)}.$$

Using (10) to rewrite the right inequality above, we obtain

$$\frac{\underline{\mu}_{\ell-1}}{1 - \underline{\mu}_{\ell-1}} > \frac{\mu \check{N} \check{y}_{\ell-1} (1 - \alpha)^{\check{N}-\ell+1}}{(1 - \mu)} \left[\frac{(1 - \alpha)^{\check{L}} (\check{y}_\ell - b_k) (\psi^\ell \delta w - 1)}{(1 - \alpha)^{\check{N}-\ell+1} \check{y}_{\ell-1} (k - 1) (\psi^{\ell-1} \delta w - 1)} \right]. \quad (\text{B.4})$$

We will show that the term in square brackets is greater than 1 and that, consequently, the expression above contradicts condition (B.3) for the $(\ell - 1)$ th in line (and for $\ell = 2$, the first inequality in (B.2)).

We begin by deriving a lower bound on $\check{y}_\ell - b_k$. If her predecessor arrived at the $(\ell - 1)$ th position in line and the next service opportunity was unsuccessful, an individual could be either the first instance of the ℓ th in line behind an uninformed first in line or any ℓ th in line behind an informed first in line. Thus,

$$b_0 + b_1 \geq \check{y}_{\ell-1}(1 - \alpha).$$

Another way to arrive at the ℓ th position in line behind an informed first in line is if the previous individual arrived at the ℓ th position and the next service opportunity produced exactly one service event:

$$b_0 \geq \check{y}_\ell \alpha (1 - \alpha) > (\check{y}_\ell - b_k) \alpha (1 - \alpha).$$

From the two inequalities above, and since the events of one service and no service are mutually exclusive, we obtain the bound

$$\check{y}_\ell - b_k \geq b_0 + b_1 \geq \check{y}_{\ell-1}(1 - \alpha) + (\check{y}_\ell - b_k) \alpha (1 - \alpha).$$

(The first inequality follows from the fact that $k \geq 2$.) Rearranging yields

$$\frac{\check{y}_\ell - b_k}{\check{y}_{\ell-1}} \geq \frac{1 - \alpha}{1 - \alpha(1 - \alpha)}.$$

Substituting, we obtain a lower bound on the terms in square brackets in (B.4):

$$\left[\frac{(1 - \alpha)^{\check{L}} (\check{y}_\ell - b_k) (\psi^\ell \delta w - 1)}{(1 - \alpha)^{\check{N} - \ell + 1} \check{y}_{\ell-1} (k - 1) (\psi^{\ell-1} \delta w - 1)} \right] \geq \underbrace{\frac{(1 - \alpha)^{\check{L} - \check{N} + \ell}}{(k - 1)}}_{K_1} \underbrace{\frac{\psi^\ell \delta w - 1}{(1 - \alpha(1 - \alpha)) (\psi^{\ell-1} \delta w - 1)}}_{K_2}.$$

We now show that both $K_1 \geq 1$ and $K_2 \geq 1$. We begin with K_1 . Observe that from (B.1), we have $k\check{L} \geq \check{N} - \ell \geq (k - 1)\check{L}$, so that

$$K_1 \geq \frac{\check{L}(1 - \alpha)^{\check{L}}}{(\check{N} - \ell)(1 - \alpha)^{\check{N} - \ell}} = \rho(1 - \alpha)^{(\rho-1)(\check{N} - \ell)} =: r(\rho),$$

where $\rho := \check{L}/(\check{N} - \ell)$. The function $r(\cdot)$ is log-concave in ρ . Moreover $r(0) = 0$ and $r(1) = 1$. Since $0 < 1/(\check{N} - \ell) < \rho \leq 1$, we have that

$$r(\rho) \geq \min \left\{ r\left(\frac{1}{\check{N} - \ell}\right), r(1) \right\}.$$

Thus, it is sufficient to show that $r(1/(\check{N} - \ell)) \geq 1$ for all $\check{N} - \ell \geq 2$ (we have already excluded the case $\check{N} - \ell = 1$). This is ensured by the condition $\alpha \geq 1/2$.

Finally,

$$K_2 \geq 1 \Leftrightarrow V_\ell > \frac{\alpha^2}{\alpha^2 - \psi(1 - \delta)} > 0.$$

In other words, $K_2 \geq 1$ is equivalent to the sufficient condition from Lemma 5 applied to the ℓ th in line. By Lemma 6, this condition is satisfied when $\delta \rightarrow 1$.

Combing all the intermediate inequalities, we have shown that the term in square brackets in (B.4) is greater than 1. Therefore, (B.4) implies

$$\frac{\underline{\mu}_{\ell-1}}{1 - \underline{\mu}_{\ell-1}} > \frac{\mu \check{N} \check{y}_{\ell-1} (1 - \alpha)^{\check{N} - \ell + 1}}{(1 - \mu)}.$$

This contradicts (B.3) for $n = \ell - 1$: the assumption that the ℓ th in line is a herding leader contradicts the assumption that agents $2, \dots, \ell - 1$ do not wish to renege before the first in line has completed her \check{N} periods of experimentation. (For $\ell = 2$, the above contradicts the first inequality in (B.2): the assumption that the second in line is a herding leader contradicts the result that the first in line experiments for \check{N} periods.) \square

Co-editor Johannes Hörner handled this manuscript.

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