Queueing to Learn*

Chiara Margaria†

October 6, 2024

Abstract

I study the efficient design of a queue to dynamically allocate a scarce resource to long-lived agents. Agents can be served multiple times, and their valuations fluctuate over time with some persistence. Each agent privately learns whether his prevailing valuation is high or low only when served. An agent can decide anytime whether to either join a queue of his choice or renege. I show that it is efficient to elicit agents' private information by offering a simple binary menu (i.e., two customer classes): a first-come, first-served queue, to attract low-value agents, and one in random order, to attract high-value agents. When queueing is costly, offering a single queue may be optimal because of the tradeoff between allocative efficiency and the cost of screening.

Keywords: Queues; Experimentation; Reneging; Congestion; Mechanism Design.
JEL Codes: C73, D47, D82

1 Introduction

This paper studies the efficient design of a queue to allocate a resource flow. Examples of rationing by waiting are plentiful: the allocation of subsidized credit and public housing, assignment of homeless shelters, provision of health and sanitation

*I am indebted to Johannes Hörner and Larry Samuelson for their invaluable help and support throughout this project. I would like to thank Dirk Bergemann, Eduardo Faingold, Mira Frick, Ryota Iijima, and Bart Lipman for their useful comments and suggestions, and Gad Allon, Barış Ata, and Achal Bassamboo for insightful discussions. I benefited from many conversations with visitors at Yale, in particular V. Bhaskar, Laura Doval, Caroline Thomas, and Takuo Sugaya, and from excellent comments from Balazs Szentes, Nikhil Vellodi, Allen I. K. Vong, and various seminar participants. I am grateful to Beixi Zhou for excellent research assistance.

†Boston University, 270 Bay State Road, Boston, MA 02215, USA, `margaria@bu.edu`.

services, allocation of donated organs, and sharing of processing power in a capacity-constrained computing system. I study the design of a queue to maximize efficiency when strategic agents require service repeatedly and learn their valuation only when served.

Consider the following stylized examples of which the model is suggestive. A microfinance institution allocates short-term loans to entrepreneurs to fund small-scale projects, such as starting a business in a developing country or increasing crop production. The profitability of each entrepreneur's project depends on market conditions and fluctuates over time. Each entrepreneur is uncertain about the prospects of his project and can assess its profitability only when allocated funds to invest; in this sense, agents learn their valuations when they are served. Because the loans are small and short-term, the same entrepreneur requests loans repeatedly to operate his business when it is profitable for him to do so.

Intuitively, the objective of the designer is twofold: maximizing allocative efficiency and minimizing congestion, that is, allocating the scarce resource to agents with the highest expected valuation (in the microfinance example, entrepreneurs with the most profitable investment opportunities) and reducing the queue length (the average time to obtain a loan).

On the face of it, a single first-come first-served queue is inefficient because agents joining the queue impose a negative externality on future arrivals by increasing the time it takes for future agents to be served, and arriving agents may have a higher expected valuation. Serving agents in a single service-in-random-order queue deters agents from joining the queue when their prevailing expected valuation is low, alleviating the externality problem. In the microfinance example, if the expected net return from a loan is negative, an entrepreneur would postpone their application until prospects improve if there is a chance of receiving the loan without delay.

More subtly, because agents are served repeatedly and learn about their valuation when they are served, the service discipline also affects the equilibrium length of the queue, or, to put it differently, the proportion of time that agents allocate to costly queueing. If loans are granted to entrepreneurs who are on average more optimistic about their return prospects, a larger fraction of them are likely to resubmit an application right away, exacerbating future congestion.

The contribution of this paper is to propose a parsimonious model to investigate these tradeoffs and to examine rationing by queueing in the absence of monetary transfers while considering the possibility of reneging.[1] The optimal queueing mechanism is remarkably simple and involves well-known queueing disciplines: it is a menu of at most two queues, service is rendered in a first-come, first-served manner in one queue and in random order in the other.

In the model, a constant flow of a resource is to be allocated to a continuum of forward-looking agents. Capacity is limited: over any interval of time, only a fixed mass of agents can be served. At each moment in time, agents decide whether and when to engage in (possibly) costly queueing to be served, and each of them can be served multiple times. Valuations fluctuate independently across agents, and each agent faces an experimentation problem because the lump-sum payoff collected at each service reveals the prevailing valuation, which can be high or low. Rivalry generates an externality problem, because agents ignore the fact that their actions affect overall congestion.

The setup presents a few methodological challenges. First, because of reneging, the revelation principle does not apply. Second, even if the underlying valuation is binary, because of learning, an agent's type belongs to a continuum. Third, the lack of transfers prevents the use of standard methods based on the envelope theorem to elicit private information.

I overcome these challenges by showing that without loss of generality, one can restrict attention to queues that provide the agents incentives not to renege. This observation, which may be useful in other contexts, relies solely on the assumption that agents acquire information about their valuation only when they are served. Further, I prove that as long as agents' expected valuations evolve monotonically over time—a condition automatically satisfied in the case of binary underlying valuations—attention can be restricted to binary menus.

I then show that as far as payoffs are concerned, each queue in the menu can be summarized by a pair of sufficient statistics. This pair determines whether a queueing discipline or menu of queueing disciplines is feasible, that is, whether the induced service rate does not exceed capacity. The idea is that both payoff and service rate depend only on the amount of time an agent expects to queue between two consecutive services and on the probability of having a high valuation when served. A version of the familiar single-crossing property of preferences holds. Specifically,

[1]Following the operations research terminology, I use the term *reneging* to describe the act of leaving a queue before being served.

the risk attitude of an agent (i.e, his preference towards a more or less risky queueing discipline) depends on his belief about his current valuation.

Having characterized the optimal menu, I then derive a set of comparative statics results. When waiting is costless or the resource is relatively abundant, it is efficient to serve agents in two queues. In contrast, when waiting is particularly wasteful, it is optimal to offer a single first-come, first-served queue, because it minimizes queue length.[2]

Related Literature. This paper is related to several strands of literature. First, it belongs to the literature on strategic behavior in queues.[3] The idea that in a first-come, first-served (FCFS) queue rational agents adopt suboptimal behavior dates back to Naor (1969). In that framework, Hassin (1985) shows that a last-come, first-served (LCFS) queueing discipline achieves the social optimum without the need for transfers (see also Scarsini and Shmaya, 2024). Platz and Østerdal (2017) find that in a concert queueing game, FCFS and LCFS achieve the minimal and maximal aggregate equilibrium payoff among all queueing disciplines, respectively. In their environment, as in mine, FCFS provides incentives to join the queue early, which in the end hurts all agents in equilibrium. The results in Hassin (1985) and Platz and Østerdal (2017), however, rely on the designer's ability to prevent restarting. In contrast, the designer in my model cannot detect or punish reneging and restarting, and LCFS cannot be part of an equilibrium.

Second, the dynamic allocation of objects to agents arriving over time through waiting lists has been studied in the context of public housing and organ transplants. Both Leshno (2022) and Bloch and Cantala (2017) consider the problem of allocating a sequence of heterogeneous items that are sequentially offered to agents on a waiting list. Both papers assume that the flow of agents joining the pool is exogenous, so maximizing welfare is equivalent to maximizing allocative efficiency; consequently, the tradeoff between allocative efficiency and congestion is absent in these models. Bloch and Cantala (2017) assume that agents' valuations evolve over time indepen-

[2]The result resonates with the anecdotal evidence that inspired Milner and Olsen (2008). The authors report that a call center was offering differentiated services to its two types of customers (those with and without a service-level agreement requiring a given percentage of customers to be served within a given time) only in off-peak hours.

[3]The motivation and modeling choices of my paper are close to those of Bassamboo and Randhawa (2015), who study scheduling policies in a queueing system with reneging customers, abstracting from strategic considerations.

dently across periods, and show that service in a first-come, first-served order always outperforms service in random order. Leshno (2022) assumes that agents' valuations are constant over time and shows that service in random order increases welfare, as compared to first-come, first-served order, by partially shielding agents from random fluctuations in waiting time. Recently, Che and Tercieux (2023) assume that the designer chooses both the queueing discipline and the information available to the agents,[4] and show that it is optimal to provide no information about queue length and serve agents according to a first-come, first-served rule.

Third, the paper contributes to the literature on dynamic mechanism design with unobservable arrival. With the exception of a few papers—among others, in the context of dynamic mechanism design with transfers, Garrett (2016) and Bergemann and Strack (2022)—most of the literature on dynamic mechanism design assumes that the designer observes agents' arrival. The case of unobservable arrival is natural in the context of queues: while the designer observes agents joining the queue, she may be unable to detect when an agent has balked and rejoined the queue, presumably disguised as a new agent. Hence, the designer is unable to condition on an agent's past history of allocation so that neither quota mechanisms as in Jackson and Sonnenschein (2007) nor a "quantified entitlement" mechanism as in Guo and Hörner (2020) is feasible. As a result, the designer elicits private information by leveraging agents' preferences about the distribution of service time, that is, their intertemporal preferences.

Last, the interaction between agents who engage in individual experimentation has been studied by the strategic experimentation literature. However, most of it has focused on information externalities, which are absent in my model. An exception is Thomas (2021), who analyzes congestion externalities. Cripps and Thomas (2019) investigate the interaction between information externalities and congestion externalities in a queueing model. Their paper is, however, only tangentially related to mine. In their model, agents arrive over time and there is a common source of uncertainty, the service rate of a server. Observational learning arises because queue length and other agents' reneging decisions reveal the agents' private information.

[4]In my model, information design would not help: because of the continuum of agents and the assumption that the queue is unobservable, agents do not need to infer the distribution of the residual waiting time from the amount they have been waiting.

The rest of the paper is organized as follows. Section 2 introduces the model. Section 3 sets up the designer's problem and simplifies it in three steps. In Section 4, I solve the designer's problem and characterize the optimal menu of queueing disciplines. Section 5 concludes the paper.

2 Setup

Time is continuous, indexed by $t \geq 0$, and the horizon is infinite. A designer (she) wishes to allocate a perishable and indivisible good to a unit mass of long-lived agents indexed by $i \in [0, 1]$. Units of the good arrive at rate λ, so a mass $\lambda (t'' - t')$ is to be assigned over any interval of time $[t', t'']$, $t' < t''$.

Let N_t^i denote the total number of times agent $i \in [0, 1]$ has been allocated the good in the time interval $[0, t]$ (formally, N_t^i is a counting process); feasibility requires that for all $t' \geq 0$ and all $t' < t''$,[5]

$$\int_0^1 \left(\int_{t'}^{t''} \mathrm{d}N_t^i \right) \mathrm{d}i \leq \lambda (t'' - t') . \tag{1}$$

The designer aims to maximize aggregate payoffs in some equilibrium of the game, as defined below.

Designer's Choice. There are no transfers, and the good is allocated via a queueing system. Informally, the designer commits to a menu of queues; at all times, each agent chooses whether to queue and which queue to join. Queues may differ in their capacity, i.e., in the amount of resource allocated to them, and in their queueing discipline. The discipline dictates how the good is allocated across the agents in the queue based on their individual waiting time. Common examples of queueing disciplines are first-come, first-served (FCSF), last-come, first-served (LCFS), and service in random order (SIRO).

A menu of queues defines an anonymous sequential game between agents (Jovanovic and Rosenthal, 1988). I focus on steady-state equilibria. Owing to the law of large numbers, in the steady state, each agent faces a single-agent decision problem.

[5]To deal with essentially pairwise independent processes $\left(N_t^i\right)_{t \geq 0}$, one needs to work with an enrichment of the usual product probability space (see Sun, 2006). While the explicit construction is omitted, I rely on Sun's law of large numbers for a continuum of random variables.

The behavior of the other agents is relevant only inasmuch as it affects the waiting time before service. Hence, I formalize the choice of the designer as a choice of a collection of (steady-state) waiting-time distributions.

More precisely, a menu is a collection of queues $\mathbf{Q}$ with generic element q. At any time t, an agent can either be queueing or not; hence, I describe agent i's action by $q_t^i \in \hat{\mathbf{Q}} := \mathbf{Q} \cup \{\varnothing\}$, with the interpretation that $q_t^i = \varnothing$ when the agent does not queue and $q_t^i \in \mathbf{Q}$ when he is waiting at some queue.

In light of the discussion above, the designer chooses a collection of (steady-state) waiting-time distributions, $\{H_q\}_{q \in \mathbf{Q}}$ for each $q \in \mathbf{Q}$, with $H_q \in \mathcal{H}$, as defined below.[6]

Definition 1. *The set $\mathcal{H}$ is the set of cumulative distribution functions, $H\colon \mathbf{R}_+ \to [0,1]$, such that $H(0) = 0$ and $\int t \,\mathrm{d}H(t) < \infty$.*

Each waiting-time distribution H_q is the distribution of the time an agent waits before being served, provided that he does not abandon the queue before service, in steady state. The requirement that H_q should not have atoms at 0 guarantees that the allocation is well defined for all strategy profiles of the agents. As will become clear (see Lemma 1), there is no loss in restricting attention to distributions having finite mean.

To understand the relationship between queueing disciplines and waiting-time distributions, note that if a queue operates in a first-come, first-served manner and agents do not renege, in steady state, all agents wait the same deterministic amount of time before being served. In other words, the waiting-time distribution is degenerate. Similarly, in a service-in-random-order queue, each agent in line, irrespectively of the amount of time he has been waiting, has the same probability of being served in the next instant. As a result, the waiting time is exponentially distributed. For ease of exposition, I restate the relationship between some waiting-time distributions and queueing disciplines in the following definitions.

Definition 2.

(i) *Agents are served according to a* first-come, first-served discipline *if the waiting-time distribution H is a degenerate distribution.*

[6] Formulating the designer's problem as a choice of waiting-time distributions circumvents the need to formalize the anonymous sequential game and the issues arising from having a continuum of independent continuous-time Markov processes. As will become clear, it amounts to restating the designer's problem as a static mechanism design problem.

(ii) *Agents are served according to a* service-in-random-order discipline *if the waiting-time distribution H is an exponential distribution with support* $[0, \infty)$.

(iii) *Agents are served according to a* service-in-random-order discipline with a minimum waiting-time requirement *if the waiting-time distribution has a constant hazard rate and support* $[t, \infty)$*, for some* $t > 0$.

Two remarks are in order. First, the designer's choice is not restricted to the classes of waiting-time distributions in the definition above, see Definition 2. Second, while the service-in-random-order discipline with a minimum waiting-time requirement is a generalization of the service-in-random-order discipline, distinguishing between them is convenient because the latter plays a more crucial role in the optimal menu characterization.

Agents' Actions. Arriving agents do not observe agents already waiting (i.e., queue length). An agent's strategy specifies when to join and leave a queue. Because time is continuous, the formal definition requires some care. In particular, an agent who leaves a queue at time t, either because he reneges or because he is served at t, may want to restart queueing with no delay.

Informally, if the agent is not queueing, a (pure) strategy dictates the time at which he joins a queue and which one. If the agent is queueing, a (pure) strategy specifies the time at which the agent reneges, i.e., leaves the queue if by that time he has not been served. In this case, the strategy prescribes the action to be taken when reneging: rejoining a queue or not. Finally, a strategy also prescribes whether to rejoin a queue with no delay after being served.

Let the time-in-queue process $(w_t^i)_{t \geq 0}$ describe the amount of time elapsed since the agent last joined the queue he is currently in whenever the agent is queueing; set the time-in-queue process equal to 0 when the agent is not queueing. An agent's strategy is an impulse control for the processes $(q_t^i)_{t \geq 0}$ and $(w_t^i)_{t \geq 0}$.[7]

If the agent does not intervene, the processes evolve exogenously as follows. While an agent is queueing, the time-in-queue grows linearly over time until the agent is served, when the time-in-queue jumps to 0. At any service time, the queue process jumps to $\varnothing$. In other words, unless the agent reneges, the agent leaves the queue as soon as he is served.

[7]The formal definition of strategies as impulse control policies is relegated to the Appendix.

At any time t, the agent can intervene and induce (q_t^i, w_t^i) to jump to either $(\varnothing, 0)$ or $(q, 0)$, for some $q \in \mathbf{Q}$. Intuitively, if a queueing agent leaves a queue—either because he reneges or because he is served—and does not rejoin the queue immediately, the process (q_t^i, w_t^i) jumps to $(\varnothing, 0)$. If the agent either joins a queue or jockeys between queues, the process (q_t^i, w_t^i) jumps to $(q, 0)$, for some $q \in \mathbf{Q}$.

Payoffs. When allocated the good, agent i receives a lump-sum payoff equal to some state θ^i, which can take two possible values, θ_0 and θ_1, with $\theta_0 < 0 < \theta_1$. Agent i's state evolves unbeknown to him according to a continuous-time Markov chain $(\theta_t^i)_{t \geq 0}$, with state space $\{\theta_0, \theta_1\}$, transition matrix $((-\rho_0, \rho_0), (\rho_1, -\rho_1))$, and initial probability of state θ_1 given by $\rho_0 / (\rho_0 + \rho_1)$. The individual state processes of any pair of agents is assumed to be independent.

Given some integrable queue process $(q_t^i)_{t \geq 0}$, the realization of the state process $(\theta_t^i)_{t \geq 0}$, and the individual allocation process $(N_t^i)_{t \geq 0}$, the realized payoff of agent i is given by the long-run average,

$$\limsup_{T \to \infty} \frac{1}{T} \left(\int_0^T \theta_t^i \, \mathrm{d}N_t^i - \int_0^T c \, \mathbf{1}_{q_t^i \neq \varnothing} \, \mathrm{d}t \right) .$$

This payoff has two components: the sum of lump-sum payoffs collected at each consumption experience and the total cost borne by the agent while queueing. Note the absence of discounting.

Strategies and Equilibrium. Given a menu, each agent faces a single-agent Markov decision problem. I introduce a state variable to describe an agent's information about his current valuation. Let p_t^i be the belief that agent i attaches to his valuation being equal to θ_1. As long as the agent is not served, his belief about this valuation evolves according to (the first-order)

$$\mathrm{d}p_t^i = \left(\left(1 - p_t^i\right) \rho_0 - p_t^i \rho_1 \right) \mathrm{d}t.$$

Specifically, for all $t' \geq 0$ and all $t' < t''$, such that no service occurs from t' to t'',

$$p_{t''}^i = e^{-(\rho_0 + \rho_1)(t'' - t')} p_{t'}^i + \left(1 - e^{-(\rho_0 + \rho_1)(t'' - t')}\right) \frac{\rho_0}{\rho_0 + \rho_1} . \tag{2}$$

Equation (2) makes it plain that the belief of agent i is a convex combination of his past belief $p^i_{t'}$ and the invariant probability of θ_1, $\rho_0/(\rho_0+\rho_1)$. Along the history with no service, the posterior belief that the state is θ_1 converges to $\rho_0/(\rho_0+\rho_1)$. As soon as the agent is served, his belief about his valuation jumps to 1 or 0.

It is without loss of generality to assume that agents' strategies are Markov in calendar time, posterior belief, current queue, and time-in-queue (t, p^i_t, q^i_t, w^i_t). Abusing notation, I denote by Σ the set of stationary Markov strategies, which are those that do not condition on calendar time.

As explained above, the agent's problem is formalized as an impulse control: the agent chooses the (random) dates at which he intervenes and adjusts his action, that is, the dates at which he joins or leaves a queue, in addition to choosing whether to rejoin a queue immediately after reneging or being served. Notice that, given an initial state (p, q, w), unless the agent adjusts his action by joining or leaving a queue, the evolution of the belief p is a sufficient statistic for the time-in-queue evolution. Hence, there is no loss of generality in focusing on impulse-control policies that are Markov in the belief.

I focus on symmetric steady-state equilibria. In a steady state, each agent $i \in [0,1]$ chooses his strategy σ^i to maximize

$$V(\sigma^i, \{H_q\}_{q\in\mathbf{Q}}) := \mathbf{E}_{\sigma^i,\{H_q\}_{q\in\mathbf{Q}}}\left[\limsup_{T\to\infty}\frac{1}{T}\left(\int_0^T \theta^i_t \,\mathrm{d}N^i_t - \int_0^T c\,\mathbf{1}_{q^i_t\neq\varnothing}\,\mathrm{d}t\right)\right]. \tag{3}$$

Definition 3. *A symmetric steady-state equilibrium is a pair* $(\sigma, \{H_q\}_{q\in\mathbf{Q}})$, $\sigma \in \Sigma$ *such that the strategy* $\sigma \in Sigma$ *is optimal given* $\{H_q\}_{q\in\mathbf{Q}} \subset \mathcal{H}$.

The designer's goal is to choose a queueing menu to maximize aggregate payoffs in some equilibrium of the game. More precisely, she chooses a symmetric steady-state equilibrium $(\sigma, \{H_q\}_{q\in\mathbf{Q}})$ to maximize aggregate payoffs. Because, by definition, the equilibrium is symmetric, each agent achieves the same realized payoff (3). Hence, the aggregate payoff equals the payoff of a representative agent, denoted by i hereafter.

The designer faces the aggregate capacity constraint (1). In the spirit of the law of large numbers, I state the capacity constraint as a bound on the average (long-run) service rate,[8]

$$\int_0^1 \left(S(\sigma^i, \{H_q\}_{q\in\mathbf{Q}})\right) \mathrm{d}i \leq \lambda, \qquad \text{where } S(\sigma^i, \{H_q\}_{q\in\mathbf{Q}}) \coloneqq \frac{1}{t} N_t^i$$

where the limit inside the integrand is understood in the sense of almost sure convergence.[9,10] For notational convenience, I drop the superscript i hereafter.

3 Designer's Problem

I simplify the designer's problem in two steps. First, I show that even if an agent's type (i.e., his belief at each point in time) belongs to a continuum, attention can be restricted to binary menus. Second, I show that the problem can be cast into a lower-dimensional space of sufficient statistics.

3.1 Binary Menus Suffice

Depending on the waiting-time distribution, an agent may benefit from leaving the queue before being served. For example, consider a waiting-time distribution with a piecewise constant hazard rate exhibiting one downward jump at some $t > 0$. If the agent were to renege and immediately rejoin the queue whenever his time-in-queue reaches t, he would effectively be served according to an exponential distribution with a rate equal to the hazard rate before t. This reasoning hinges on the fact that the designer is unable to detect restarting (the combined action of reneging and rejoining). Even if reneging is optimal, any best reply satisfies a few desirable properties.

Proposition 1. *If $\sigma \in \Sigma$ maximizes $V(\sigma, \{H_q\}_{q\in\mathbf{Q}})$, and $V(\sigma, \{H_q\}_{q\in\mathbf{Q}}) > 0$, then, on path, if $q_t = \varnothing$ and $w_{t-} > 0$, $p_t = 0$. That is, the agent stops queueing only when he realizes a lump sum θ_0.*

[8] Note that the law of the counting process N_t^i is jointly determined by $\{H_q\}_{q\in\mathbf{Q}}$ and σ^i, but for notational simplicity, I keep such dependence implicit.

[9] In Section A.2.1, I show that for any $\{H_q\}_{q\in\mathbf{Q}}$ and any best reply σ, the long-run service rate converges to a constant almost surely.

[10] As I will show, the solution to the designer's problem involves waiting-time distributions, which are easily implementable with well-known queueing disciplines. As a result, it is not necessary to argue that for any equilibrium $(\sigma, \{H_q\}_{q\in\mathbf{Q}}) \in \Sigma \times \mathcal{H}^{\mathbf{Q}}$ such that $S(\sigma, \{H_q\}_{q\in\mathbf{Q}}) \leq \lambda$, it is possible to find a collection of queueing disciplines implementing it. This is a collection of allocation rules such that the induced anonymous sequential game between agents has a symmetric equilibrium in which each player adopts the strategy σ and the collection of waiting-time distributions is $\{H_q\}_{q\in\mathbf{Q}}$.

To put it differently, regardless of whether and when an agent reneges, starting from the time when he first joins a queue, he queues uninterruptedly until he realizes a lump sum θ_0, possibly jockeying between different queues in the intervening time. Roughly, there are two reasons why an agent may join a queue: he may want to be served as soon as possible or at some point in the future. In the first case, it is intuitive that reneging and spending any amount of time balking before rejoining the queue is suboptimal. In the second case, if the agent found it optimal to renege and wait before rejoining the queue, it would not have been optimal for him to join the queue so early.

Proposition 1 has two implications. First, when best-replying, if the agent reneges, he rejoins the queue with no delay. Second, the agent rejoins the queue immediately after realizing a high lump-sum payoff. As a result, a menu of queueing disciplines can be understood as a set of possible induced distributions. We can distinguish two types of agents ("low type" and "high type"), depending on the realized payoff at the last service (θ_0 and θ_1) and hence on whether their belief is below or above the invariant probability of state θ_1, $\rho_0/(\rho_0+\rho_1)$. Faced with a menu, each type of agent selects the preferred waiting-time distribution from those that can be "engineered" by repeatedly leaving and joining (potentially) different queues over time.

When the agent plays the strategy σ, the distribution of the amount of time he spends in queue between two consecutive services is either $\hat{H}_0^\sigma$ or $\hat{H}_1^\sigma$, depending on the realized payoff at the last service. To formally define these distributions, let $\{T_n\}_{n=1,2,\dots} := \{t \in [0,\infty) \mid \mathrm{d}N_t > 0\}$ be the sequence of service times. The two induced distributions are:

$$\begin{aligned} \hat{H}_0^\sigma(t) &:= \Pr_{\sigma,\{H_q\}_{q\in\mathbf{Q}}}\left[T_{n+1}-\tau \le t \mid T_n \le \tau < T_{n+1}^i, p_\tau = \underline{p}^\sigma\right], \\ \hat{H}_1^\sigma(t) &:= \Pr_{\sigma,\{H_q\}_{q\in\mathbf{Q}}}\left[T_{n+1}-T_n \le t \mid p_{T_n} = 1\right], \end{aligned} \tag{4}$$

where $\underline{p}^\sigma$ is the lowest belief at which the agent joins the queue when playing σ, $\underline{p}^\sigma := \liminf_{t\to\infty}\{p_t : q_t \neq \varnothing\}$.

Combining these observations, it follows that it is without loss of generality to restrict attention to binary menus and non-reneging strategies. The next lemma, the proof of which is in the spirit of the revelation principle, states this formally. Let $\Sigma^{\mathrm{NR}} \subset \Sigma$ be the set of non-reneging strategies.

Lemma 1. *If $\sigma \in \Sigma$ is optimal given $\{H_q\}_{q\in\mathbf{Q}}$, then there exists a binary menu $\{H_0, H_1\} \subset \mathcal{H}$ and a strategy $\sigma' \in \Sigma^{\mathrm{NR}}$ such that*

(i) *σ' is optimal within Σ given $\{H_0, H_1\}$;*

(ii) *under σ', for any $t \geq 0$, $q_t = 1$ if and only if $p_t = 1$; and*

(iii) *$(\sigma', \{H_0, H_1\})$ and $(\sigma, \{H_q\}_{q\in\mathbf{Q}})$ yield the same payoffs and the same service rate, i.e., $V(\sigma, \{H_q\}_{q\in\mathbf{Q}}) = V(\sigma', \{H_0, H_1\})$ and $S(\sigma, \{H_q\}_{q\in\mathbf{Q}}) = S(\sigma', \{H_0, H_1\})$.*

Proof. Set $H_0 = \hat{H}_0^\sigma$, and $H_1 = \hat{H}_1^\sigma$. If when offered the menu $\{\hat{H}_0^\sigma, \hat{H}_1^\sigma\}$, the agent uses the best non-reneging strategy that prescribes choosing $\hat{H}_1^\sigma$ at check-in when his belief of his valuation is 1 and $\hat{H}_0^\sigma$ otherwise, his payoff is unchanged. Moreover, any other strategy is suboptimal. In fact, when faced with the original menu, the agent is able to induce the waiting-time distributions $\hat{H}_0^\sigma$ and $\hat{H}_1^\sigma$ and hence any possible convolution of truncations thereof, but he finds it optimal to induce the distributions $\hat{H}_0^\sigma$ and $\hat{H}_1^\sigma$. Hence, he has a best reply $\sigma' \in \Sigma^{\mathrm{NR}}$ satisfying (ii). □

Property (ii) is an incentive-compatibility condition: when checking in, each agent finds it optimal to join the queue designed for his type. Given Lemma 1, the designer's problem can be restated as

$$\text{(M)} \qquad \sup V(\sigma, \{H_0, H_1\})$$

over binary menus $\{H_0, H_1\} \subset \mathcal{H}$ and strategies $\sigma \in \Sigma^{\mathrm{NR}}$, subject to

$$\sigma \in \arg\max_{\sigma^i \in \Sigma} V(\sigma^i, \{H_0, H_1\}), \tag{5}$$

$$S(\sigma, \{H_0, H_1\}) \leq \lambda, \tag{C}$$

$$(\hat{H}_0^\sigma, \hat{H}_1^\sigma) = (H_0, H_1). \tag{IC}$$

Restricting the designer to non-reneging strategies is without loss of generality. In fact, the designer's problem is reminiscent of a static delegation problem and could be stated as a choice between pairs of induced waiting-time distributions. The chosen pair would then correspond to an equivalence class of strategies: the constraint (IC) provides the criterion for selecting one strategy from this class. Additionally, even if the designer is constrained to non-reneging strategies, each agent, when best-replying, is not restricted to this class, as is clear from the incentive compatibility constraint (5).

3.2 Non-Reneging Constraint

It might appear that requiring the designer to select $\sigma \in \Sigma^{\mathrm{NR}}$ makes the incentive problem more difficult: she must provide incentives to play threshold strategies, and agents' incentives to do so depend on fine details of the distributions H_0 and H_1. However, the following proposition shows that it is sufficient to restrict attention to the class of new-better-than-used-in-expectation (NBUE) distributions, which I define next.

Definition 4. *A distribution $H \in \mathcal{H}$ is called NBUE (new better than used in expectation) if for all $t > 0$,*

$$\frac{\int_t^\infty (1 - H(s))\,\mathrm{d}s}{1 - H(t)} \leq \int_0^\infty (1 - H(s))\,\mathrm{d}s. \tag{6}$$

Let $\mathcal{H}^{\mathrm{NBUE}} \subset \mathcal{H}$ denote the set of NBUE distributions.[11] In reliability theory, NBUE distributions are used to describe the lifetime of a non-repairable component.[12] In the context of queues, if a queue is characterized by an NBUE waiting-time distribution, an agent in that queue expects to wait an amount of time (the left-hand side of (6)) no longer than the expected waiting time of an agent who has just joined the queue (the right-hand side of (6)). Intuitively, restarting the queue cannot reduce the residual expected waiting time when the waiting-time distribution belongs to the NBUE class.

Clearly, degenerate distributions and exponential distributions belong to the NBUE class. In fact, when the discipline is first-come, first-served, restarting the queue strictly increases the residual expected waiting time, whereas the residual expected waiting time is constant if agents are served in random order. Additionally, it is a well-known result in reliability theory that any NBUE distribution is less variable than an exponential distribution with the same mean in the following sense: a waiting-time distribution is NBUE if and only if it second-order stochastically dominates an exponential distribution with the same mean (Ross, 1995, Proposition 9.6.1).

Proposition 2.

[11] Strictly speaking, the concept of NBUE (see Shaked and Shanthikumar, 2007) is defined for any nonnegative distribution with finite mean, while $\mathcal{H}$ excludes distributions with atoms at 0.

[12] In economics, Gershkov and Moldovanu (2010) (see also Chapter 2 in Gershkov and Moldovanu, 2014) leverage the properties of the class of NBUE distributions to bound the efficient policy in a sequential allocation problem with incomplete information.

(i) *Restricting attention to NBUE distributions, it is without loss to assume that the agent is restricted to the class of non-reneging strategies. That is, if* $\sigma \in \Sigma^{\text{NR}}$ *is optimal within* Σ^{NR} *given* $\{H_0, H_1\} \subset \mathcal{H}^{\text{NBUE}}$*, then there exists* $\{H_0', H_1'\} \subset \mathcal{H}^{\text{NBUE}}$ *such that*

  (a) σ *is optimal within* Σ *given* $\{H_0', H_1'\}$*;*

  (b) $(\sigma, \{H_0', H_1'\})$ *and* $(\sigma, \{H_0, H_1\})$ *induce the same payoffs and the same service rate.*

(ii) (a) *A binary menu is incentive compatible only if the distribution* H_1 *is NBUE. That is, if* $\sigma \in \Sigma^{\text{NR}}$ *is optimal within* Σ *given* $\{H_0, H_1\} \subset \mathcal{H}$ *and satisfies* (IC)*, then* $H_1 \subset \mathcal{H}^{\text{NBUE}}$*.*

  (b) *A binary menu is optimal only if the distribution* H_0 *is NBUE. That is, if* $\{H_0, H_1\} \subset \mathcal{H}$ *together with* $\sigma \in \Sigma^{\text{NR}}$ *solves* (M)*, then* $H_0 \subset \mathcal{H}^{\text{NBUE}}$*.*

The first part of the proposition states that when focusing on the class of NBUE distribution, the designer does not need to worry about reneging. For any menu $\{H_0, H_1\} \subset \mathcal{H}^{\text{NBUE}}$ and any non-reneging strategy $\sigma \in \Sigma^{\text{NR}}$, one can find a pair of "payoff and service rate equivalent" NBUE distributions $\{H_0', H_1'\} \subset \mathcal{H}^{\text{NBUE}}$ that provide incentives not to renege.

The second part shows that without loss of optimality, the designer can restrict attention to the class of NBUE distributions. First, as explained above, an NBUE waiting-time distribution has the property that an agent, when queueing, faces an expected residual waiting time no longer than the average waiting time of newcomers. As a consequence, agents who are becoming pessimistic about their individual state find it optimal to restart the queue whenever the waiting-time distribution does not satisfy the NBUE property and, $H_1 \in \mathcal{H}^{\text{NBUE}}$ is a necessary condition for agents to have an incentive not to renege.

One cannot use the same argument for agents who are becoming optimistic about their individual valuation, as they may not want to be served as soon as possible. As I shall explain in Section 4, "low types" dislike randomness in their service time, while the opposite is true for "high types". Intuitively, in light of the characterization of NBUE distributions in terms of second-order stochastic dominance and the fact that $H_1 \in \mathcal{H}^{\text{NBUE}}$, it follows that choosing a distribution $H_0 \notin \mathcal{H}^{\text{NBUE}}$ cannot be beneficial for screening purposes and is suboptimal.

3.3 Finite-Dimensional Problem

Even restricting attention to NBUE distribution, solving for the optimal binary menu of waiting-time distributions remains an infinite-dimensional problem. In this section, I show that given a menu, the payoff and the service rate from any stationary Markov strategy can be written as a function of a few sufficient statistics, and the designer's problem can be cast in the finite-dimensional space of sufficient statistics.

As discussed in Section 2, there is no loss of generality in focusing on impulse-control policies that are Markov in the belief. Strictly speaking, this description of strategies (in terms of the posterior belief only) does not specify the behavior of an agent following off-path histories.[13] However, with respect to payoffs, this is innocuous in my environment. Even if the agent were to start the game with an arbitrary posterior belief, the transient component of the payoff (that is, the payoff collected before reaching the recurrent path) would not affect the (long-run average) payoff.[14]

Sufficient Statistics. Fix a best reply $\sigma \in \Sigma$. The payoff process induced by σ can be expressed as a function of the belief process. This follows from two observations. First, as discussed above, without loss, the action is only a function of the belief, and so is the flow cost incurred by the agent. Second, the times at which the belief jumps coincide with the times at which the agent collects lump sums. Because when the agent is served, the belief takes extreme values, to compute the payoff and the service rate from the strategy σ, we can focus on the two-state Markov renewal process with state space $\{0, 1\}$ one obtains by sampling the belief process at the service times, as illustrated in Figure 1.

We can use this two-state representation to compute the payoff by assuming that at each visit to state $s = 0, 1$, the agent collects a reward equal to $\theta_s - c\,\hat{\mu}_{\sigma,s}$ where $\hat{\mu}^\sigma_s$,

$$\hat{\mu}^\sigma_s := \mathbf{E}_{\sigma,\{H_q\}_{q\in\mathbf{Q}}}\left[\int_{T_n}^{T_{n+1}} \mathbf{1}_{q_t \neq \varnothing}\,\mathrm{d}t \mid p_{T_n} = s\right], \qquad s = 0, 1$$

[13] That is, the strategy does not specify the behavior after one's own deviation.

[14] There is a small caveat: there exist strategies such that given an initial belief p_0, the recurrent path is never reached. I discuss these "absorbing" strategies at the end of the proof of Lemma 2, in the Appendix. Absorbing strategies cannot be part of an equilibrium, and hereafter, I focus on non-absorbing strategies.

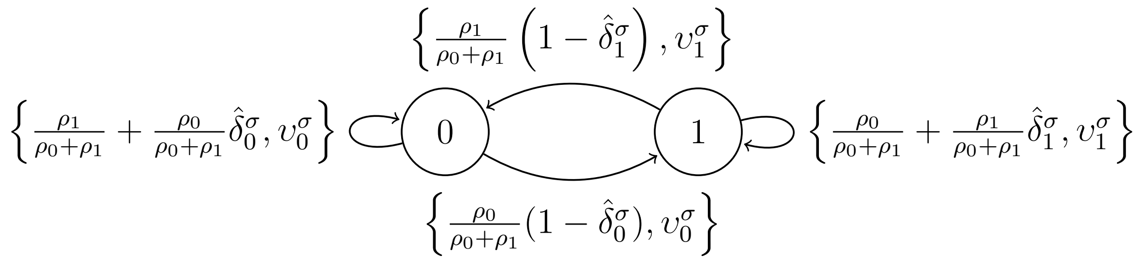

Figure 1: Markov Renewal Process. The expression in brackets represents the transition probability and the average sojourn time. The computation of these variables can be found in the proof of Lemma 2.

is the expected time that an agent spends in queue before the next service, starting from the moment the belief reaches s.[15]

By the definition of induced waiting-time distribution, (4), the expected time the agent spends in queue before the next service starting from the moment the belief reaches 0 and 1 is, respectively,

$$\hat{\mu}_0^\sigma = \int t \, \mathrm{d}\hat{H}_0^\sigma(t), \qquad \hat{\mu}_1^\sigma = \int t \, \mathrm{d}\hat{H}_1^\sigma(t).$$

The sojourn times and the transition probabilities of the Markov renewal process are determined by the strategy σ and the menu of waiting-time distributions via two pairs of statistics, defined next. Define, for $s = 0, 1$,

$$\begin{aligned} \upsilon_s^\sigma &:= \mathbf{E}_{\sigma, \{H_q\}_{q \in \mathbf{Q}}} \left[T_{n+1} - T_n \mid p_{T_n} = s \right], \\ \hat{\delta}_s^\sigma, &:= \mathbf{E}_{\sigma, \{H_q\}_{q \in \mathbf{Q}}} \left[e^{-(\rho_0+\rho_1)(T_{n+1}-T_n)} \mid p_{T_n} = s \right]. \end{aligned}$$

The first statistic, υ_s^σ, is the expected amount of time between two consecutive services. Because time enters the evolution of the posterior belief of an agent exponentially, see (2), it is not surprising that the second statistic, $\hat{\delta}_\sigma^\sigma$, is the moment generation function evaluated at $-(\rho_0+\rho_1)$. The probability that the Markov renewal process transitions from state s to state 1 is equal to the probability that, after being served at time T_n, at the next service, which occurs at T_{n+1}, the agent realizes a

[15]Recall that $\{T_n\}_{n=1,2,\ldots} := \{t \in [0, \infty) \mid \mathrm{d}N_t > 0\}$ denote the service times.

payoff of θ_1, that is

$$\mathbf{E}_{\sigma,\{H_q\}_{q\in\mathbf{Q}}}\left[e^{-(\rho_0+\rho_1)(T_{n+1}-T_n)} \mid p_{T_n}=s\right] s \\ + \left(1-\mathbf{E}_{\sigma,\{H_q\}_{q\in\mathbf{Q}}}\left[e^{-(\rho_0+\rho_1)(T_{n+1}-T_n)} \mid p_{T_n}=s\right]\right)\frac{\rho_0}{\rho_0+\rho_1}. \tag{7}$$

The weight attached to the initial belief in the convex combination is the expectation of the exponential function that appears in (2) and, by definition, is equal to $\hat{\delta}_s^\sigma$.

The following identities make it plain that the statistics depend on the strategy σ and the menu of waiting-time distributions only through the induced waiting-time distributions and the lowest belief at which the agent join the queue:

$$\hat{\delta}_1^\sigma = \int e^{-(\rho_0+\rho_1)t}\,\mathrm{d}\hat{H}_1^\sigma(t), \qquad \hat{\delta}_0^\sigma = \varrho\left(\underline{p}^\sigma\right)\int e^{-(\rho_0+\rho_1)t}\,\mathrm{d}\hat{H}_0^\sigma(t),$$
$$\upsilon_1^\sigma = \hat{\mu}_1^\sigma, \qquad \upsilon_0^\sigma = -\ln\varrho\left(\underline{p}^\sigma\right)/(\rho_0+\rho_1)+\hat{\mu}_0^\sigma,$$

where

$$\varrho(p) = \left|1-\frac{p}{\rho_0/(\rho_0+\rho_1)}\right|, \tag{8}$$

is a function of the time required for the belief to increase from 0 to p.

Once the transition and reward structure of the Markov renewal process induced by a strategy σ are defined, the computation of payoffs is standard. Lemma 2 summarizes the result. The proof is in the Appendix.

Lemma 2. *Fix $\{H_q\}_{q\in\mathbf{Q}}\subset\mathcal{H}$, and $\sigma\in\Sigma$. Then, $V(\sigma,\{H_q\}_{q\in\mathbf{Q}})$ and $S(\sigma,\{H_q\}_{q\in\mathbf{Q}})$ are a function of υ_s^σ, $\hat{\delta}_{\sigma,s}^\sigma$, $\hat{\mu}_s^\sigma$, $s=0,1$, only.*

Feasible Statistics. As shown above, the relevant statistics are a function of the induced waiting-time distribution, that ultimately will be equal to the offered waiting-time distribution. Hence, it is convenient to define, for any distribution H, the following summary statistics:

$$\mu^H := \int_0^\infty t\,\mathrm{d}H(t)\,, \qquad \delta^H := \int_0^\infty e^{-(\rho_0+\rho_1)t}\,\mathrm{d}H(t)\,.$$

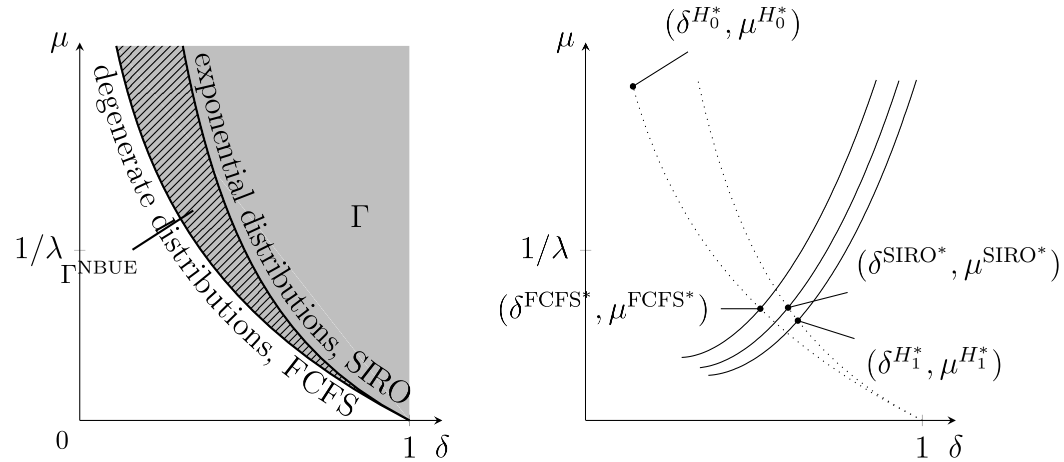

Figure 2: On the left panel, the sets Γ and Γ^{NBUE}. On the right panel, summary statistics for the optimal menu $\{H_0^*, H_1^*\}$ and for the best disciplines within the SIRO and FCFS classes. Solid lines depict the agent's indifference curves; utility is increasing in the southeast direction. $(\theta_1, \theta_0, c, \rho_0, \rho_1, \lambda) = (1, -3/4, 0, 1, 1, 2)$.

To identify for which pairs (δ, μ) there exists a distribution $H \in \mathcal{H}$ ($H \in \mathcal{H}^{\text{NBUE}}$) such that $(\delta, \mu) = \left(\delta^H, \mu^H\right)$, I use Lemma 3. Its proof, statistical in nature, is relegated to the Appendix.

Lemma 3. *The following are equivalent:*

(i) *There exists a distribution $H \in \mathcal{H}$ ($H \in \mathcal{H}^{\text{NBUE}}$) such that $\delta^H = \delta$ and $\mu^H = \mu$;*

(ii) *it holds that $(\delta, \mu) \in \Gamma$ ($(\delta, \mu) \in \Gamma^{\text{NBUE}}$), where*

$$\Gamma := \left\{(\delta, \mu) \in (0,1) \times (0, \infty) \mid e^{-(\rho_0+\rho_1)\mu} \leq \delta \leq 1\right\},$$
$$\Gamma^{\text{NBUE}} := \left\{(\delta, \mu) \in (0,1) \times (0, \infty) \mid e^{-(\rho_0+\rho_1)\mu} \leq \delta \leq \mathbf{E}\left[e^{-(\rho_0+\rho_1)\text{Exp}(\mu)}\right]\right\},$$

and $\text{Exp}(\mu)$ *is an exponential random variable with mean* μ.

The sets Γ and Γ^{NBUE} are depicted in the left panel of Figure 2. Recall that for any waiting-time distribution H, the two statistics (δ^H, μ^H) are, respectively, the moment generating function evaluated at $-(\rho_0 + \rho_1)$ and the expected waiting time. Because the function $t \mapsto e^{-(\rho_0+\rho_1)t}$ is convex, the southwestern boundary of the sets corresponds to degenerate distributions, that is, those assigning probability 1 to some $\mu \in (0, \infty)$. The northeastern boundary of Γ^{NBUE} corresponds to the set of exponential distributions, which are, as mentioned above, "extreme" within the NBUE family. From the perspective of incentives, when the waiting time is exponentially distributed, the non-reneging constraint is binding at all times: agents are served at a constant rate, independent of their arrival time in the queue.

A noteworthy consequence of Lemma 3 is that the classes of waiting-time distributions corresponding to the three classes of queueing disciplines in Definition 2 span the set Γ^{NBUE}. If agents are served according to a first-come, first-served discipline, the pair of summary statistics lies on the western boundary of Γ^{NBUE}. If agents are served according to a service-in-random-order discipline, the pair of summary statistics lies on the eastern boundary of Γ^{NBUE}. Finally, each point in the interior of Γ^{NBUE} is achieved by a shifted exponential distribution that can be generated by serving agents in random order with a minimum waiting-time requirement $t > 0$.

Notice, however, that the characterization in Lemma 3 does not account for the capacity constraint. On the one hand, if the designer offers a single queue and does not discard any of the available resource flow, the expected waiting time does not exceed $1/\lambda$. On the other hand, as I shall explain in the next section, identifying the best feasible first-come-first-served queue, for example, is not merely a statistical problem as it requires analyzing the agent's best reply.

4 Optimal Menu

Since the designer maximizes the aggregate payoff and each agent achieves the same payoff in equilibrium, the designer's preferences coincide with each agent's preferences. However, there is scope for the intervention by a designer because agents do not internalize the externality generated by their actions. To shed light on the problem faced by the designer, I now present a payoff decomposition that highlights the source of the externality. (The formal derivation can be found in the Appendix.)

Fix a pair of waiting-time distributions $\{H_0, H_1\} \subset \mathcal{H}$ and a strategy $\sigma \in \Sigma^{NR}$ that satisfies (IC). The payoff can be written as

$$V(\sigma^i, \{H_0, H_1\}) = S(\sigma^i, \{H_0, H_1\}) \cdot \Big[m(\delta^{H_0}, \delta^{H_1}, \underline{p}^{\sigma}) \left(\theta_1 - c\,\mu^{H_1}\right) + \left(1 - m(\delta^{H_0}, \delta^{H_1}, \underline{p}^{\sigma})\right) \left(\theta_0 - c\,\mu^{H_0}\right) \Big], \tag{9}$$

where $m(\delta^{H_0}, \delta^{H_1}, \underline{p}^{\sigma})$ is the long-run frequency with which a service yields a lump sum θ_1, and

$$S(\sigma^i, \{H_0, H_1\}) = \frac{1}{m(\delta^{H_0}, \delta^{H_1}, \underline{p}^{\sigma})\mu_1 + \left(1 - m(\delta^{H_0}, \delta^{H_1}, \underline{p}^{\sigma})\right)\left(\mu_0 - \ln \varrho\left(\underline{p}^{\sigma}\right)/(\rho_0 + \rho_1)\right)}$$

is the induced service rate.

According to (9), the payoff from the strategy σ is the product of the rate at which the agent is served and the expected total payoff he collects between service times. The latter is a function of $m(\delta^{H_0}, \delta^{H_1}, \underline{p}^{\sigma})$, the probability of being served when the state is θ_1. The relationship between $m(\delta^{H_0}, \delta^{H_1}, \underline{p}^{\sigma})$ and $S(\sigma^i, \{H_0, H_1\})$ is easy to understand. From the elementary renewal theorem, the expected service rate equals the inverse of the average time between services. When joining the queue, the agent expects to wait an amount of time equal to either μ^{H_0} or μ^{H_1}, depending on the payoff he realized at the last service. Moreover, after being served, he waits an amount of time $-\ln \varrho\left(\underline{p}^{\sigma}\right)/(\rho_0 + \rho_1)$ before joining the queue whenever the realized payoff is θ_0, which occurs a proportion $1 - m(\delta^{H_0}, \delta^{H_1}, \underline{p}^{\sigma})$ of the time.

The cutoff $\underline{p}^{\sigma}$ affects the rate at which an agent is served. This is a manifestation of the congestion externality, reminiscent of a "tragedy of the commons". The designer must guarantee through an appropriate choice of distributions that the service rate induced by the agents' best reply does not exceed the capacity λ. Ideally, the designer would like to minimize wasteful wait and persuade the agents with a low belief to delay as long as possible before joining the queue.

The cutoff $\underline{p}^{\sigma}$ also affects the rate of arrival of the high types, or to put it differently, the representative agent's probability of realizing a high lump-sum payoff when served. The longer the agent waits before joining the queue after realizing a lump sum payoff of θ_0, the larger rate of arrival of the high types. At the same time, by Proposition 1, an agent rejoins the queue with no delay as soon as he collects a high lump-sum payoff, so a larger rate of arrival of the high types may increase the aggregate queueing cost.

4.1 A Special Case: Costless Queueing

I start by characterizing the optimal menu when queueing is costless, to highlight the tradeoff stemming from the dynamic externality problem. When queueing is costless, the negative externality that an agent imposes on another agent is not related to the cost of queuing but rather to the rivalry between agents. Agents' desire to be served can be due to an exploration or exploitation motive: agents who want to explore, because they are growing increasingly optimistic about their valuation, do not internalize the fact that they may curtail other agents' ability to exploit, that is, to be served when their expected valuation is the highest.

To develop some intuition regarding the optimal menu, notice that when queueing is costless, serving agents in a single service-in-random-order queue yields a higher payoff compared to serving them in a single first-come, first-served queue. When agents are served in random order, they may be served immediately after realizing a high lump-sum payoff when their expected valuation is the highest, which never happens when serving them in order of arrival. The right panel of Figure 2 plots the summary statistics for the best feasible service-in-random-order queue and the best feasible first-come, first-served queue: serving agents in random order may involve a longer wait compared to serving them in order of arrival, but because queueing is costless, the former is welfare-improving.

Now, observe that the designer could achieve the same payoff and service rate as a service in random order queue while having agents queue at all times by offering a binary menu consisting of a service-in-random-order queue and a service-in-random-order queue with a minimum waiting-time requirement. To put it differently, when $c = 0$, the designer does not need to try and persuade agents with a low belief to delay joining the queue, and without loss of generality, we can assume that in the optimal menu, agents queue at all times.

Next, I argue that a version of the familiar single-crossing property of preferences holds: from the law of motion of beliefs (7), an agent's attitude toward uncertainty in the service time—whether he is risk-seeking or risk-averse over time lotteries—depends on whether he is growing optimistic or pessimistic about his valuation. An agent with a belief below the invariant probability dislikes randomness in his service time, while the opposite is true for agents with a belief above the invariant probability.

Naturally, one of the incentive constraints must bind; for otherwise, the designer could decrease the wait at the service-in-random-order queue and increase the wait at the first-come, first-served queue and increase payoffs, without violating any constraint. In the optimal menu, the incentive constraint of the agent joining with a low belief binds. As a result, the optimal menu is payoff-equivalent to serving agents in a single service-in-random-order queue. Of course, to achieve the same payoff with a single service-in-random-order queue, the designer would need a larger capacity than λ: the wait at the service-in-random-order queue in the menu is shorter than the wait at the best feasible service-in-random-order queue (see the right panel of Figure 2).

Theorem 1. *Suppose $c = 0$. An optimal menu is one such that $\mu^{H_1^*} \geq \mu^{H_0^*}$, $\delta^{H_0^*} \geq \delta^{H_1^*}$, the capacity constraint* (C) *is binding, and*

(i) *(FCFS/SIRO menu) H_0^* is degenerate and H_1^* is exponential;*

(ii) *(low-type IC binds) any best reply to $\{H_1^*, H_1^*\}$ yields the same payoff as $(\sigma, \{H_0^*, H_1^*\})$; and*

(iii) *(agents queue at all times) $\underline{p}^\sigma = 0$.*

Agents are offered a choice between two queues: one with a first-come, first-served discipline and the other with a random-order discipline. The agents joining the queue with a high belief, that is, immediately after receiving a positive lump-sum payoff, are served in random order, the "riskiest" discipline that provides incentives not to renege. The agents joining immediately after receiving a negative lump-sum payoff are served according to a first-come, first-served queueing discipline; hence, they are exposed to minimal risk.

4.2 The General Case

When $c > 0$, considerations about queue length cannot be ignored. The trade-off between allocative efficiency and congestion is more subtle, and pooling different types of agents by offering a single queue is sometimes optimal.

Theorem 2. *There exists a solution $(\sigma, \{H_0^*, H_1^*\})$ to the designer's problem* (M). *It is such that $\mu^{H_1^*} \geq \mu^{H_1^*}$, $\delta^{H_1^*} \geq \delta^{H_1^*}$, the capacity constraint* (C) *is binding, and one of the following holds:*

(i) *(pooling menu) $H_0^* = H_1^* = H^*$ for some $H^* \in \mathcal{H}^{\text{NBUE}}$,*

(ii) *(separating menu) $H_0^* \neq H_1^*$, and*

   (a) *(FCFS/SIRO menu) H_0^* is degenerate and H_1^* is exponential;*

   (b) *(low-type IC binds) any best reply to $\{H_1^*, H_1^*\}$ yields the same payoff as $(\sigma, \{H_0^*, H_1^*\})$; and*

   (c) *(agents queue at all times) $\underline{p}^\sigma = 0$.*

The optimal menu can be of two types: pooling or separating. Intuitively, in the absence of monetary transfers, queueing is not only a byproduct of scarcity but also serves as a costly signaling device. A screening menu requires agents to engage in wasteful queueing to signal their type and allocates dedicated capacity to the "high types." As in Condorelli (2012), sometimes the designer finds it optimal not to extract agents' private information and instead offers a single queue.

When the optimal menu is pooling, the optimal service discipline is either first-come, first-served or service in random order with or without a minimum waiting-time requirement. As shown below, both first-come, first-served and service-in-random-order (with an arbitrary waiting-time requirement $t > 0$) disciplines can emerge as optimal for some sets of parameters $(\theta_1, \theta_0, \rho_0, \rho_1, \lambda)$.

A separating optimal menu coincides with the one in Theorem 1. The value of the information acquired at each service is maximized: learning is so valuable that agents queue at all times, even if queueing is costly. From the perspective of the individual experimentation problem, this does not mean that exploring is valuable at every belief. Joining the queue has an option value: it guarantees the right to be served at some point in the future when the belief will be higher. An agent joins the queue at a belief of 0 because he is certain that he will not engage in exploration for some time.

4.3 Discussion: Comparative Statics

As mentioned before, when $c > 0$, considerations about queue length cannot be ignored. In fact, the separating menu, if optimal, maximizes queue length. To the other extreme, if the waiting cost is high enough,[16] the designer finds it optimal to offer a single first-come, first-served queue, even if this implies forgoing the possibility to serve returning agents as soon as they rejoin the queue.

Lemma 4. *Fix any admissible set of parameters $(\theta_1, \theta_0, \rho_0, \rho_1, \lambda)$.*

(i) *There exists a $\overline{c}$ such that for $c > \overline{c}$, neither the separating menu nor the pooling service-in-random order queue is optimal.*

[16] Because rescaling (θ_1, θ_0, c) amounts to rescaling payoffs but does not affect the implementable set Γ, increasing c is equivalent to decreasing the gain from targeting the high types $\theta_1 - \theta_0$.

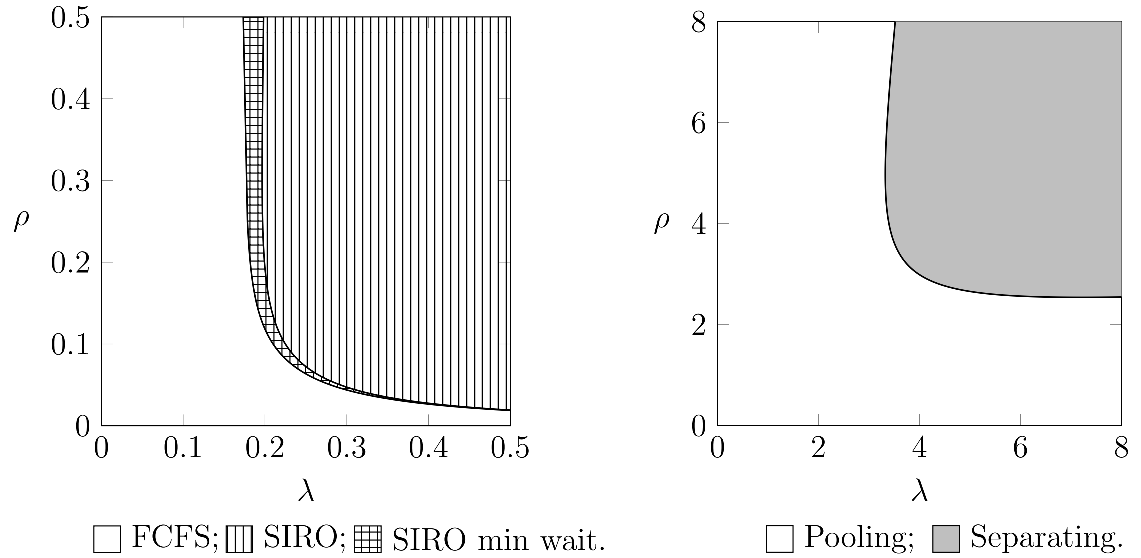

Figure 3: Comparative statics for $(\theta_1, \theta_0, c) = (1, -3/4, 1/4)$ and $\rho_0 = \rho_1 = \rho$. The shading indicates the features of the optimal queueing discipline.

(ii) *If offering a single first-come, first-served queue is optimal, then it minimizes the average waiting time μ (equivalently, the queue length) among all feasible disciplines.*

Part (i) formalizes the idea that the benefit from serving agents who are likely to have a high prevailing valuation does not pay off for the increased congestion when the queueing cost is high, while part (ii) lays bare the fact that the benefit from a first-come, first-served queueing discipline comes from shortening the queue.

The comparative statics with respect to λ are summarized in Lemma 5 and Figure 3. When the resource is scarce, it is optimal to offer a single queue. A first-come, first-served queue is suboptimal for a high enough λ.

Lemma 5. *Fix any admissible set of parameters $(\theta_1, \theta_0, \rho_0, \rho_1)$.*

(i) *There exists a $\underline{\lambda}$ such that for $\lambda < \underline{\lambda}$, the separating menu is not optimal.*

(ii) *There exists a $\overline{\lambda}$ such that offering a single first-come, first-served queue is suboptimal for any $\lambda \geq \overline{\lambda}$.*

The role of persistence is less clear-cut. On the one hand, the informational value from being served increases with persistence, making screening more valuable. On the other hand, the cost of not serving returning agents as soon as they join the queue decreases with persistence. In the extreme case, if the state becomes arbitrarily persistent, all non-reneging single-queue service disciplines perform equally well. The ambiguous impact of persistence is shown in Figure 3, where for simplicity, I set $\rho_0 = \rho_1 = \rho$. As shown in the right panel, for some set of parameters, offering two queues is optimal only for an intermediate range of the persistence parameter ρ. While analytical results are difficult to obtain, numeral simulations suggest that the patterns identified in Figure 3 are the rule.

Although the solution may be different for different objective functions, the analysis provides the tools to study the queueing discipline that minimizes waiting time or maximizes the expected return from service.

In the working paper (Margaria, 2021), I show that the optimal menu can be virtually implemented, in the sense that it is possible to implement an outcome arbitrarily close to it with a single queueing discipline by taking advantage of reneging.

5 Concluding Remarks

I study the optimal design of a queue to allocate a resource to agents with heterogeneous preferences. In this setup, a menu to screen agents takes the form of multiple queues (or customer classes), with agents being served in a different order within each of them. The optimal menu is (at most) binary and has a simple structure. When it is optimal to offer two distinct queues, agents are served in a first-come, first-served manner in one queue and in random order in the other queue. When pooling is optimal, the single queue is either first-come, first-served or service in random order, possibly with a minimum waiting-time requirement.

The analysis rests on three main assumptions: anonymity, no transfers, and learning. It can be shown that if any of the first two assumptions is dropped, it is possible to restore the first best, that is, to achieve a total payoff arbitrarily close to $\lambda\theta_1$.[17] If agents observe the evolution of their state, the ranking of queueing disciplines is unambiguous: the service-in-random-order discipline dominates any other queueing discipline. If the designer is able to detect reneging, the optimal menu may involve a discipline that exposes the high-valuation agents to maximal variability in waiting time by serving at each point in time some of the newly arrived (as in a last-come,

[17] Each of these extensions is examined in the working paper, Margaria (2021).

first-served queue) and some of the agents who have been waiting the longest (as in a first-come, first-served queue).

The model is stylized in many respects. Some assumptions are for convenience. For instance, the optimal menu is binary even if the valuation takes more than two values, provided that the expected valuation $\mathbf{E}\left[\theta_t^i \mid \theta_t^i = \theta_j\right]$ evolves monotonically over time. The assumption of a continuum of agents makes it possible to formulate the best-reply problem as a simple Markov decision problem. A model with a finite number of agents would allow for a finer analysis of the strategic interaction between them, beyond the general-equilibrium effect captured by the current model. However, to the extent that the problem reduces to a "two-level" optimization problem, I believe that the main insights would not be overturned in a setting with a large but finite number of agents.

Allowing for variable capacity would be useful to study the welfare implications of congestion in an environment with fluctuations.[18] More broadly, because of the stationarity of the environment and the independence assumption, important aspects of queueing and learning via experimentation are missing from the current model.[19] For example, correlation in agents' valuations would introduce the possibility of observational learning.

[18] Interestingly, the peer-to-peer lending platform Zopa, a two-tier queueing mechanism to allocate lenders' funds to borrowing opportunities, prioritizes returning lenders. However, one may expect economic fluctuations to play a key role in this environment.

[19] Similarly, the assumption of a deterministic capacity, allows to abstract away from exogenous randomness in the service time, as would be the case if units of the good arrived stochastically.

A Proofs

A.1 Preliminaries

This section contains the formal definition of strategies as impulse control policies. Fix a menu $\{H_q : q \in \mathbf{Q}\}$ and let $(\mathcal{F}_t^i)_{t\geq 0}$ be the filtration corresponding to the information of agent i. A strategy for agent i, $i \in [0,1]$, is a sequence of random times and random variables, (intervention times and impulses at these times, respectively),

$$\sigma = \{(\tau_k, \varsigma_k)\}_{k=1}^{\infty},$$

where

(i) $0 \leq \tau_1 < \tau_2 < \dots,$

(ii) for any $k \in \mathbf{N}$,

$$\tau_k = \min\left\{\left\{t > \tau_{k-1} : \mathrm{d}N_t^i > 0\right\}, \tilde{\tau}_k\right\},$$

where $\tilde{\tau}_k$ is a predictable stopping time adapted to the filtration $(\mathcal{F}_t^i)_{t\geq 0}$ (predictability enforces the informational restriction that, when queueing at t, an agent chooses the stopping time τ at which he reneges conditional on the event $\{\inf\{t' > t : \mathrm{d}N_{t'}^i > 0\} \geq \tau\}$),

(iii) for any $k \in \mathbf{N}$, $\varsigma_k \in \hat{\mathbf{Q}}$ is a $\mathcal{F}_{\tau_k}^i$-measurable random variables, and

(iv) for any $k \in \mathbf{N}$, if $\varsigma_k = \varnothing$, then $\varsigma_{k+1} \neq \varnothing$ a.s.

The strategy defines the action process $(q_t^i, w_t^i)_{t\geq 0}$ taking values in $\hat{\mathbf{Q}} \times \mathbf{R}$, where $(q_t^i)_{t\geq 0}$ is piecewise constant, $q_{\tau_k}^i = \varsigma_k$, and $w_t^i := (t - \sup_{\tau_k \leq t} \tau_k)\mathbf{1}_{q_t^i \neq \varnothing}$.

A.2 Proofs for Section 3

I first show that the payoff and the service rate from any stationary Markov strategy can be written as a function of a few sufficient statistics and prove the characterization of feasible statistics in Lemma 3, as stated in Section 3.3. Then, I prove the results in Section 3.1 and Section 3.2.

A.2.1 Proof of Lemma 2

The transition matrix of the semi-Markov chain in Figure 1 is

$$\begin{pmatrix} \frac{\rho_1}{\rho_0+\rho_1} + \frac{\rho_0}{\rho_0+\rho_1}\hat{\delta}_0^\sigma & \frac{\rho_0}{\rho_0+\rho_1}(1-\hat{\delta}_0^\sigma) \\ \frac{\rho_1}{\rho_0+\rho_1}(1-\hat{\delta}_1^\sigma) & \frac{\rho_1}{\rho_0+\rho_1} + \frac{\rho_0}{\rho_0+\rho_1}\hat{\delta}_1^\sigma \end{pmatrix}.$$

The chain is positive recurrent and irreducible. The unique stationary distribution is

$$\begin{pmatrix} 1-m \\ m \end{pmatrix} := \begin{pmatrix} \dfrac{\frac{\rho_1}{\rho_0+\rho_1}(1-\hat{\delta}_1^\sigma)}{\frac{\rho_1}{\rho_0+\rho_1}(1-\hat{\delta}_1^\sigma) + \frac{\rho_0}{\rho_0+\rho_1}(1-\hat{\delta}_0^\sigma)} \\ \dfrac{\frac{\rho_0}{\rho_0+\rho_1}(1-\hat{\delta}_0^\sigma)}{\frac{\rho_1}{\rho_0+\rho_1}(1-\hat{\delta}_1^\sigma) + \frac{\rho_0}{\rho_0+\rho_1}(1-\hat{\delta}_0^\sigma)} \end{pmatrix}.$$

First, assume $\nu_s^\sigma < \infty$, $s = 0, 1$. Because any $H \in \mathcal{H}$ has no atoms at 0, there exist $\varepsilon > 0$ and $\varepsilon' > 0$, such that $\Pr\left[T_n - T_{n-1} \leq \varepsilon\right] \leq 1-\varepsilon$. Moreover the transition matrix is unichain. Consequently, the long-run average payoff can be computed using the evaluation equations (by Puterman, 1994, Th. 11.4.2, Ch. 11). The long-run average payoff equals

$$\frac{m\,(\theta_1 - c\,\hat{\mu}_1^\sigma) + (1-m)\,(\theta_0 - c\,\hat{\mu}_0^\sigma)}{m\,\nu_1^\sigma + (1-m)\,\nu_0^\sigma} = \frac{\frac{\rho_0}{\rho_0+\rho_1}(1-\hat{\delta}_0^\sigma)\,(\theta_1 - c\,\hat{\mu}_1^\sigma) + \frac{\rho_1}{\rho_0+\rho_1}(1-\hat{\delta}_1^\sigma)\,(\theta_0 - c\,\hat{\mu}_0^\sigma)}{\frac{\rho_0}{\rho_0+\rho_1}(1-\hat{\delta}_0^\sigma)\,\nu_1^\sigma + \frac{\rho_1}{\rho_0+\rho_1}(1-\hat{\delta}_1^\sigma)\,\nu_0^\sigma}. \quad (10)$$

The time between T_n and T_{n+1}, for some $n \in \mathbf{N}$, is independent of n. Hence, the average number of upward jumps per unit of time converges almost surely to the inverse of the mean inter-arrival time (see for example Asmussen, 2003, Chapter V, Proposition 1.4), that is, the service rate converges almost surely to a constant. From direct inspection of (10), the rate at which the agent collects lump sums converges almost surely to

$$\lim_{t\to\infty} \frac{1}{t} N_t = \frac{1}{m\,\nu_1^\sigma + (1-m)\,\nu_0^\sigma} = \frac{\frac{\rho_1}{\rho_0+\rho_1}(1-\hat{\delta}_1^\sigma) + \frac{\rho_0}{\rho_0+\rho_1}(1-\hat{\delta}_0^\sigma)}{\frac{\rho_0}{\rho_0+\rho_1}(1-\hat{\delta}_0^\sigma)\,\nu_1^\sigma + \frac{\rho_1}{\rho_0+\rho_1}(1-\hat{\delta}_1^\sigma)\,\nu_0^\sigma}. \quad (11)$$

Because the game is symmetric and I analyze the steady state, the long-run fraction of time for which the agent is queueing coincides with the length of the queue in steady state and is equal to

$$\frac{m\,\hat{\mu}_{\sigma,1} + (1-m)\,\hat{\mu}_0^\sigma}{m\,\nu_1^\sigma + (1-m)\,\nu_0^\sigma} = \frac{\frac{\rho_0}{\rho_0+\rho_1}(1-\hat{\delta}_0^\sigma)\hat{\mu}_1^\sigma + \frac{\rho_1}{\rho_0+\rho_1}(1-\hat{\delta}_1^\sigma)\hat{\mu}_0^\sigma}{\frac{\rho_0}{\rho_0+\rho_1}(1-\hat{\delta}_0^\sigma)\,\nu_1^\sigma + \frac{\rho_1}{\rho_0+\rho_1}(1-\hat{\delta}_1^\sigma)\,\nu_0^\sigma}.$$

Second, assume $\nu_s^\sigma = \infty$ for some $s \in \{0, 1\}$. In this case, the expected long-run average payoff is either $-\infty$ or 0, depending on

$$\lim_{T\to\infty} \mathbf{E}_{\sigma,\{H_q\}_{q\in\mathbf{Q}}} \left[\int_0^T \mathbf{1}_{q_t\neq\varnothing}\,\mathrm{d}t\right] \gtreqless 0.$$

If the previous limit is 0, the long-run average payoff is 0; it diverges to $-\infty$ otherwise. In both cases, the service rate converges almost surely to zero. □

A.2.2 Proof of Lemma 3

Notice that the function $x \mapsto e^{-(\rho_0+\rho_1)x}$ is convex. Thus, given a mean $\mu > 0$, the minimum value for the other statistics is achieved by the random variable that is degenerate at μ. This, together with $e^{-(\rho_0+\rho_1)x} < 1$, proves that for any $H \in \mathcal{H}$, $(\delta, \mu) \in \Gamma$. For the other direction, let $(\delta, \mu) \in \Gamma$, $\delta \neq e^{-(\rho_0+\rho_1)\mu}$. Consider a distribution H that randomizes between $\{\varepsilon/\pi, (\mu-\varepsilon)/(1-\pi)\}$, with probability $(\pi, 1-\pi)$, where $0 < \varepsilon < \mu$ and $\pi > 0$ are chosen to satisfy

$$\pi\, e^{-(\rho_0+\rho_1)\varepsilon/\pi} + (1-\pi)\, e^{-(\rho_0+\rho_1)(\mu-\varepsilon)/(1-\pi)} = \delta.$$

If ε was 0, the previous equation would have a unique root $\pi \in (0, 1)$. Because the left-hand side is continuous in ε, there exist an $\varepsilon > 0$ and a $\pi \in (0, 1)$ such that the equality is satisfied. Clearly, $H \in \mathcal{H}$, and by construction, $(\delta^H, \mu^H) = (\delta, \mu)$. If instead $\delta = e^{-(\rho_0+\rho_1)\mu}$, the statistics of the random variable degenerate at μ are (δ, μ).

To show the other equivalence, let $H \in \mathcal{H}^{\mathrm{NBUE}}$. Any NBUE random variable with mean μ is smaller than $\mathrm{Exp}[\mu]$ in the convex stochastic order,[20] where $\mathrm{Exp}[\mu]$ is the

[20] The random variable X is said to be smaller than Y in the convex order if

$$\mathbf{E}\left[\phi(X)\right] \leq \mathbf{E}\left[\phi(Y)\right] \text{ for all convex functions } \phi\colon \mathbb{R} \to \mathbb{R},$$

provided the expectation exists.

exponential random variable with the mean μ (see Shaked and Shanthikumar, 2007, Chap. 3, Th. A.55). Because the function $x \mapsto e^{-(\rho_0+\rho_1)x}$ is convex, it follows that any $H \in \mathcal{H}^{\text{NBUE}}$ satisfies

$$\delta^H \leq \mathbf{E}\left[e^{-(\rho_0+\rho_1)\text{Exp}(\mu)}\right] = \frac{1}{1+(\rho_0+\rho_1)\,\mu}.$$

Consequently, any $H \in \mathcal{H}^{\text{NBUE}}$, $\left(\delta^H, \mu^H\right) \in \Gamma^{\text{NBUE}}$. To show the converse, let $(\delta, \mu) \in \Gamma^{\text{NBUE}}$. Note that any degenerate distribution belongs to $\mathcal{H}^{\text{NBUE}}$. For $\mu > 0$, let $\text{D}(\mu)$ denote the random variable degenerate at μ. Then, by the properties of the moment generating function,

$$\mathbf{E}\left[e^{-(\rho_0+\rho_1)(\alpha\text{D}(\mu)+(1-\alpha)\text{Exp}(\mu))}\right] = e^{-(\rho_0+\rho_1)\alpha\mu}\frac{1}{1+(\rho_0+\rho_1)\,(1-\alpha)\,\mu}.$$

Let α be such that

$$e^{-(\rho_0+\rho_1)\alpha\mu}\frac{1}{1+(\rho_0+\rho_1)\,(1-\alpha)\,\mu} = \delta. \tag{12}$$

Because $\delta \in \left[e^{-(\rho_0+\rho_1)\mu}, \frac{1}{1+(\rho_0+\rho_1)\mu}\right]$ and the left-hand side is decreasing in α, for $\alpha \in [0,1]$, there exists a unique root α to (12) in $[0,1]$. This shows that one can find a random variable H that is a convolution of a degenerate distribution and an exponential distribution such that $(\delta, \mu) = \left(\delta^H, \mu^H\right)$. Because convolutions of IHR distributions are IHR, the random variable $\alpha\text{D}(\mu) + (1-\alpha)\text{Exp}(\mu)$ is IHR, and hence it is an NBUE random variable. Because $\mu < \infty$, and neither $\text{D}(\mu)$, nor $0 < \text{Exp}(\mu)$ have atoms at 0, $H \in \mathcal{H}^{\text{NBUE}}$.

□

A.2.3 Proof of Proposition 1

In the proof, I assume that $\{H_q\}_{q\in\mathbf{Q}}$ is such that the agent has a best reply that yields strictly positive payoffs. (In light of Lemma 11, this assumption is without loss of generality.) As argued, a strategy $\sigma \in \Sigma$ can be described in terms of one state variable only, the posterior belief. I now introduce some notation to describe stationary Markov strategies in a way that exploits this recursivity. Fix a (pure) strategy $\sigma \in \Sigma$. For any $p \in [0,1]$, define $q_\sigma : [0,1] \to \hat{\mathbf{Q}}$ to be such that, along the path induced by σ, $q_\sigma(p_t) = q_t$ a.s. Additionally, define $\tau_\sigma \colon [0,1] \to \mathbf{R}_+$ so that, a.s, along the path induced by σ,

$$\tau_\sigma(p) := \left\{ \inf_{\tau \geq 0} \mid w_{t+\tau} - w_{t+\tau-} \neq 0 \text{ or } q_{t+\tau} \neq q_{t-} \mid p_t = p, \mathrm{d}N_{t+\tau} = 0 \right\}.$$

The maps τ_σ and q_σ completely characterize the strategy $\sigma \in \Sigma$: for any starting belief p and action (on the recurrent path induced by σ), the agent adjusts his action either after an interval of time $\tau_\sigma(p)$, or when his belief jumps, whichever occurs first.

Note that, the time it takes for the belief to go from p to p', in the absence of jumps, whenever either $p < p' < \rho/(\rho_0 + \rho_1)$ or $p > p' > \rho/(\rho_0 + \rho_1)$ is

$$\left(\ln \varrho(p) - \ln \varrho(p')\right) / (\rho_0 + \rho_1)$$

where the function $\varrho : [0, 1] \to \mathbf{R}$ was defined in (8).

Overview of the proof. The proof proceeds by contradiction and is divided into two steps. Assume that $\sigma \in \Sigma$ is optimal given $\{H_q\}_{q \in \mathbf{Q}}$. First, I show that for any $p \in [\rho_0/(\rho_0 + \rho_1)), 1]$ on the recurrent path induced by σ, $q_\sigma(p) \neq \varnothing$, that is, the agent rejoins the queue immediately after realizing a high lump-sum payoff and if he reneges, he immediately restarts the queue. Then, I show that $q_\sigma(p) \neq \varnothing$ whenever $p \in [p^0_\sigma, \rho_0/(\rho_0 + \rho_1))$.

No abandoning at $p \geq \rho_0/(\rho_0 + \rho_1)$. Assume by contradiction that $q_\sigma(p) = \varnothing$ for some $p \in [\rho_0/(\rho_0 + \rho_1), 1)$ on the recurrent path. Let $p' := \sup \{p \mid q_\sigma(p) = \varnothing\}$ and $p'' := \sup \{p < p' \mid q_\sigma(p) \neq \varnothing\}$. Because the best reply yields a strictly positive payoff, $\rho_0/(\rho_0 + \rho_1) < p'' \leq p' \leq 1$, and by the initial assumption, $p'' < p'$.

Define the strategy $\tilde{\sigma} \in \Sigma$ such that

$$\tau_{\tilde{\sigma}}(p) = \begin{cases} \tau_\sigma(p) & p < \rho_0/(\rho_0 + \rho_1), \\ \tau_\sigma \left(\frac{\varrho(p)}{\varrho(p')} \left(p'' - \frac{\rho_0}{\rho_0 + \rho_1} \right) + \frac{\rho_0}{\rho_0 + \rho_1} \right) & \rho_0/(\rho_0 + \rho_1) \leq p \leq p', \\ \tau_\sigma(p) & p' < p, \end{cases}$$

$$q_{\tilde{\sigma}}(p) = \begin{cases} q_\sigma(p) & p < \rho_0/(\rho_0 + \rho_1), \\ q_\sigma \left(\frac{\varrho(p)}{\varrho(p')} \left(p'' - \frac{\rho_0}{\rho_0 + \rho_1} \right) + \frac{\rho_0}{\rho_0 + \rho_1} \right) & \rho_0/(\rho_0 + \rho_1) \leq p \leq p', \\ q_\sigma(p) & p' < p. \end{cases}$$

According to the strategy $\tilde{\sigma}$, the agents rejoins the queue at p', and his behavior from then on coincides with the behavior of an agent who joined at p'' and followed the strategy σ.

From (10), the payoff is strictly decreasing in ν_1^σ and $\hat{\mu}_1^\sigma$ and strictly increasing in $\hat{\delta}_1^\sigma$. Clearly, $\nu_0^{\tilde{\sigma}} = \nu_0^\sigma$, $\hat{\mu}_0^{\tilde{\sigma}} = \hat{\mu}_0^\sigma$, and $\hat{\delta}_0^\sigma = \hat{\delta}_0^{\tilde{\sigma}}$. Notice that $\nu_1^{\tilde{\sigma}} = \nu_1^\sigma - (\ln \varrho(p') - \ln \varrho(p''))/(\rho_0 + \rho_1)$ and $\hat{\mu}_1^{\tilde{\sigma}} = \hat{\mu}_1^\sigma$. Also, because along the path induced by σ, p' is visited with positive probability, it can be checked that $\hat{\delta}_1^\sigma < \hat{\delta}_1^{\tilde{\sigma}}$ contradicting the optimality of σ.

No abandoning at $p < \rho_0/(\rho_0+\rho_1)$**.** Towards a contradiction, assume that $q_\sigma(p) = \varnothing$ for some $p > \underline{p}^\sigma$ which is reached with positive probability on the induced path. From (10), the payoff is strictly decreasing in ν_0^σ, $\hat{\mu}_0^\sigma$, and $\hat{\delta}_0^\sigma$. Let $p' \coloneqq \inf\{p > \underline{p}^\sigma \mid q^\sigma(p) = \varnothing\}$ and $p'' \coloneqq \inf\{p > p' \mid q^\sigma(p) \neq \varnothing\}$. By the initial assumption, $p' < p''$, and because the best reply yields a strictly positive payoff, $p'' < \rho_0/(\rho_0 + \rho_1)$.

I now consider a class of strategies that includes σ and show that σ is suboptimal within this class. Specifically, each strategy in this class is characterized by a pair of beliefs, $(\underline{p}^\dagger, \tilde{p}'')$ such that (i) $0 \leq \underline{p}^\dagger < \tilde{p}'' < \rho_0/(\rho_0 + \rho_1)$, and (ii) $\frac{\varrho(p')}{\varrho(\underline{p}^\sigma)}\left(\underline{p}^\dagger - \frac{\rho_0}{\rho_0+\rho_1}\right) + \frac{\rho_0}{\rho_0+\rho_1} \leq \tilde{p}''$. For a pair of beliefs $\left(\underline{p}^\dagger, \tilde{p}''\right)$ satisfying these restrictions, let $\sigma_{\underline{p}^\dagger,\tilde{p}''}$ be the strategy such that

$$\tau_{\sigma_{\underline{p}^\dagger,\tilde{p}''}}(p) = \begin{cases} \left(\ln \varrho(p) - \ln \varrho(\underline{p}^\dagger)\right)/(\rho_0 + \rho_1) & p < \underline{p}^\dagger \\ \tau_\sigma\left(\frac{\varrho(p)}{\varrho(\underline{p}^\dagger)}\left(\underline{p}^\sigma - \frac{\rho_0}{\rho_0+\rho_1}\right) + \frac{\rho_0}{\rho_0+\rho_1}\right) & \underline{p}^\dagger \leq p \leq \frac{\varrho(p')}{\varrho(\underline{p}^\sigma)}\left(\underline{p}^\dagger - \frac{\rho_0}{\rho_0+\rho_1}\right) + \frac{\rho_0}{\rho_0+\rho_1}, \\ \left(\ln \varrho(p) - \ln \varrho(\tilde{p}'')\right)/(\rho_0 + \rho_1) & \frac{\varrho(p')}{\varrho(\underline{p}^\sigma)}\left(\underline{p}^\dagger - \frac{\rho_0}{\rho_0+\rho_1}\right) + \frac{\rho_0}{\rho_0+\rho_1} < p < \tilde{p}'', \\ \tau_\sigma\left(\frac{\varrho(p)}{\varrho(\tilde{p}'')}\left(p'' - \frac{\rho_0}{\rho_0+\rho_1}\right) + \frac{\rho_0}{\rho_0+\rho_1}\right) & \tilde{p}'' \leq p \leq \rho_0/(\rho_0 + \rho_1), \\ \tau_\sigma(p) & \rho_0/(\rho_0 + \rho_1) < p, \end{cases}$$

$$q_{\sigma_{\underline{p}^\dagger,\tilde{p}''}}(p) = \begin{cases} \varnothing & p < \underline{p}^\dagger \text{ and } \frac{\varrho(p')}{\varrho(\underline{p}^\sigma)}\left(\underline{p}^\dagger - \frac{\rho_0}{\rho_0+\rho_1}\right) + \frac{\rho_0}{\rho_0+\rho_1} < p < \tilde{p}'', \\ q_\sigma\left(\frac{\varrho(p)}{\varrho(\underline{p}^\dagger)}\left(\underline{p}^\sigma - \frac{\rho_0}{\rho_0+\rho_1}\right) + \frac{\rho_0}{\rho_0+\rho_1}\right) & \underline{p}^\dagger \leq p \leq \frac{\varrho(p')}{\varrho(\underline{p}^\sigma)}\left(\underline{p}^\dagger - \frac{\rho_0}{\rho_0+\rho_1}\right) + \frac{\rho_0}{\rho_0+\rho_1}, \\ q_\sigma\left(\frac{\varrho(p)}{\varrho(\tilde{p}'')}\left(p'' - \frac{\rho_0}{\rho_0+\rho_1}\right) + \frac{\rho_0}{\rho_0+\rho_1}\right) & \tilde{p}'' \leq p \leq \rho_0/(\rho_0 + \rho_1) \\ q_\sigma(p) & \rho_0/(\rho_0 + \rho_1) < p. \end{cases}$$

In words, according to the strategy $\sigma_{\underline{p}^\dagger,\tilde{p}''}$, starting from a belief of 0, the agent joins the queue as soon as the belief reaches $\underline{p}^\dagger$. After joining, the agent behaves "as if" he had joined at $\underline{p}^\sigma$. Because σ involves reneging at p', according to $\sigma_{\underline{p}^\dagger,\tilde{p}''}$, the agent reneges at $\left(\varrho(p')/\varrho(\underline{p}^\sigma)\right)\left(\underline{p}^\dagger - \rho_0/(\rho_0+\rho_1)\right) + \rho_0/(\rho_0+\rho_1)$. Note that this is not necessarily the first time the agent reneges after joining: σ may involve jockeying between queues multiple times, as would $\sigma_{\underline{p}^\dagger,\tilde{p}''}$. The second parameter $\tilde{p}''$ is the belief at which the agent rejoins: according to $\sigma_{\underline{p}^\dagger,\tilde{p}''}$, the agent rejoins at $\tilde{p}''$ and thereafter follows the path of actions prescribed by σ after p''.

By construction, $\hat{\mu}_0^{\sigma_{\underline{p}^\dagger,\tilde{p}''}} = \hat{\mu}_0^\sigma$, $\hat{\mu}_1^{\sigma_{\underline{p}^\dagger,\tilde{p}''}} = \hat{\mu}_1^\sigma$, $\nu_1^{\sigma_{\underline{p}^\dagger,\tilde{p}''}} = \nu_1^\sigma$ and

$$\begin{aligned}\nu_0^{\sigma_{\underline{p}^\dagger,\tilde{p}''}} =&\Pr_{\sigma,\{H_q\}_{q\in\mathbf{Q}}}\left[T_{n+1} < -\ln\left(\varrho(p')\right)/(\rho_0+\rho_1) \mid p_{T_n}=0\right] \\ &\Bigg(\mathbf{E}_{\sigma,\{H_q\}_{q\in\mathbf{Q}}}\left[T_{n+1}-T_n \mid p_{T_n}=0, T_{n+1} < -\ln\left(\varrho(p')\right)/(\rho_0+\rho_1)\right] \\ &\quad + \left(\ln\varrho(\underline{p}^\sigma) - \ln\varrho(\underline{p}^\dagger)\right)/(\rho_0+\rho_1)\Bigg) \\ &+ \Pr_{\sigma,\{H_q\}_{q\in\mathbf{Q}}}\left[T_{n+1} \geq -\ln\left(\varrho(p')\right)/(\rho_0+\rho_1) \mid p_{T_n}=0\right] \\ &\Bigg(\mathbf{E}_{\sigma,\{H_q\}_{q\in\mathbf{Q}}}\left[T_{n+1}-T_n \mid p_{T_n}=0, T_{n+1} \geq -\ln\left(\varrho(p')\right)/(\rho_0+\rho_1)\right] \\ &\quad + \left(\ln\varrho(p'') - \ln\varrho(\tilde{p}'')\right)/(\rho_0+\rho_1)\Bigg),\end{aligned} \tag{13}$$

and

$$\begin{aligned}&\hat{\delta}_0^{\sigma_{\underline{p}^\dagger,\tilde{p}''}} \\ &= \Pr_{\sigma,\{H_q\}_{q\in\mathbf{Q}}}\left[T_{n+1} < -\ln\left(\varrho(p')\right)/(\rho_0+\rho_1) \mid p_{T_n}=0\right]\frac{\varrho\left(\underline{p}^\dagger\right)}{\varrho\left(\underline{p}^\sigma\right)} \\ &\quad\cdot \mathbf{E}_{\sigma,\{H_q\}_{q\in\mathbf{Q}}}\left[e^{-(\rho_0+\rho_1)(T_{n+1}-T_n)} \mid p_{T_n}=0, T_{n+1} < -\ln\left(\varrho(p')\right)/(\rho_0+\rho_1)\right] \\ &+ \Pr_{\sigma,\{H_q\}_{q\in\mathbf{Q}}}\left[T_{n+1} \geq -\ln\left(\varrho(p')\right)/(\rho_0+\rho_1) \mid p_{T_n}=0\right]\frac{\varrho\left(\tilde{p}''\right)}{\varrho\left(p''\right)} \\ &\quad\cdot \mathbf{E}_{\sigma,\{H_q\}_{q\in\mathbf{Q}}}\left[e^{-(\rho_0+\rho_1)(T_{n+1}-T_n)} \mid p_{T_n}=0, T_{n+1} \geq -\ln\left(\varrho(p')\right)/(\rho_0+\rho_1)\right].\end{aligned} \tag{14}$$

The choice of the optimal strategy within this class can be formulated as a choice over pairs $\left(\varrho\left(\underline{p}^\dagger\right)/\varrho\left(\underline{p}^\sigma\right), \varrho\left(\tilde{p}''\right)/\varrho\left(p''\right)\right)$ subject to suitable feasibility constraints.

Specifically, constraint (i) and (ii) imply that the following constraint must hold:

$$\begin{cases} \dfrac{\varrho\left(\underline{p}^{\dagger}\right)}{\varrho\left(\underline{p}^{\sigma}\right)} \leq \dfrac{\varrho\left(0\right)}{\varrho\left(\underline{p}^{\sigma}\right)} \\ \dfrac{\varrho\left(\tilde{p}''\right)}{\varrho\left(p''\right)} \leq \dfrac{\varrho\left(p'\right)}{\varrho\left(p''\right)} \dfrac{\varrho\left(\underline{p}^{\dagger}\right)}{\varrho\left(\underline{p}^{\sigma}\right)} \end{cases} \tag{15}$$

I conclude the proof by showing that the optimal pair must satisfy (15) with equality. Proceeding by contradiction, suppose that the optimal pair is interior. Then, using the fact that the payoff is strictly monotonic in the sufficient statistics,

$$\frac{\dfrac{\partial \nu_0^{\sigma_{\underline{p}^{\dagger},\tilde{p}''}}}{\partial\left(\varrho\left(\underline{p}^{\dagger}\right)/\varrho\left(\underline{p}^{\sigma}\right)\right)}}{\dfrac{\partial \nu_0^{\sigma_{\underline{p}^{\dagger},\tilde{p}''}}}{\partial\left(\varrho\left(\tilde{p}''\right)/\varrho\left(p''\right)\right)}} = \frac{\dfrac{\partial \hat{\delta}_0^{\sigma_{\underline{p}^{\dagger},\tilde{p}''}}}{\partial\left(\varrho\left(\underline{p}^{\dagger}\right)/\varrho\left(\underline{p}^{\sigma}\right)\right)}}{\dfrac{\partial \hat{\delta}_0^{\sigma_{\underline{p}^{\dagger},\tilde{p}''}}}{\partial\left(\varrho\left(\tilde{p}''\right)/\varrho\left(p''\right)\right)}}. \tag{16}$$

However, from (13) and (14), this is impossible, because

$$\frac{\frac{\varrho\left(\underline{p}^{\dagger}\right)}{\varrho\left(\underline{p}^{\sigma}\right)}\mathbf{E}_{\sigma,\{H_q\}_{q\in\mathbf{Q}}}\left[e^{-(\rho_0+\rho_1)(T_{n+1}-T_n)} \mid p_{T_n}=0, T_{n+1} < -\ln\left(\varrho(p')\right)/(\rho_0+\rho_1)\right]}{\frac{\varrho(\tilde{p}'')}{\varrho(p'')}\mathbf{E}_{\sigma,\{H_q\}_{q\in\mathbf{Q}}}\left[e^{-(\rho_0+\rho_1)(T_{n+1}-T_n)} \mid p_{T_n}=0, T_{n+1} \geq -\ln\left(\varrho(p')\right)/(\rho_0+\rho_1)\right]} > 1,$$

whereas (16) holds true only if this ratio equals 1. Because the constraint (15) is slack when evaluated at σ (as $p' < p''$), this concludes the proof that any optimal strategy must satisfy the properties in Proposition 1.

A.2.4 Proof of Proposition 2

Proof of (i)

If $H_0 \times H_1 \in \mathcal{H}^{\text{NBUE}} \times \mathcal{H}^{\text{NBUE}}$, by Lemma 3, $(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) \in \Gamma^{\text{NBUE}} \times \Gamma^{\text{NBUE}}$. As shown in the proof of Lemma 3, there exists a pair of convolutions of an exponential distribution and a degenerate distribution H_0' and H_{θ_1}' such that $(\delta^{H_{\theta_0}'}, \mu^{H_0'}) = (\delta^{H_0}, \mu^{H_0})$ and $(\delta^{H_{\theta_1}'}, \mu^{H_1'}) = (\delta^{H_1}, \mu^{H_1})$. In other words, H_{θ_0}' and H_{θ_1}' are either degenerate distributions or shifted (to the right) exponential distributions. Clearly, in the first case, agents have no incentives to renege, as they are served in order of arrival. In the second case, reneging and rejoining is dominated because it introduces "gaps" in the induced waiting-time distribution (see the proof of Proposition 1). Consequently, the optimal strategy within Σ given $\{H'_{\theta_0}, H'_{\theta_1}\}$ is a non-reneging strategy. In addition, since the optimal non-reneging strategy is a function of the two pairs of summary statistics only, σ is optimal given $\{H_0, H_1\}$. Similarly, by the results in Section 3.3, the service rate is unchanged because it is a function of σ and summary statistics only.

Proof of (ii)(a).

I first state and prove a lemma which is used later.

Lemma 6. *If* (H_0, H_1) *is incentive compatible and* H_1 *has unbounded support, then*

$$\limsup_{t\to\infty} \frac{\int_t^\infty 1 - H_1(s)\,\mathrm{d}s}{1 - H_1(t)} \leq \int_0^\infty 1 - H_1(t)\,\mathrm{d}t. \tag{17}$$

Proof. Assume by contradiction, that

$$\limsup_{t\to\infty} \frac{\int_t^\infty 1 - H_1(s)\,\mathrm{d}s}{1 - H_1(t)} - \int_0^\infty 1 - H_1(t)\,\mathrm{d}t > 0.$$

For $\{t_n\}_{n\geq 1}$, $t_n \to \infty$ consider the strategy that prescribes joining the queue H_1 after realizing a high-lump sum payoff and restarting it (once only) at t_n. The induced changed in the two statistics are, respectively,

$$\begin{aligned}\int_0^{t_n} e^{-(\rho_0+\rho_1)s}\,\mathrm{d}H_1(s) + (1 - H_1(t_n))e^{-(\rho_0+\rho_1)t_n}\left(\int_0^\infty e^{-(\rho_0+\rho_1)(s)}\,\mathrm{d}H_1(s)\right)&\\ - \int_0^\infty e^{-(\rho_0+\rho_1)s}\,\mathrm{d}H_1(s) = o(1 - H_1(t_n))&\\ 0 > H_1(t_n)\int_0^{t_n} 1 - \frac{H_1(s)}{H_1(t_n)}\,\mathrm{d}s + (1 - H_1(t_n))\left(t_n + \int_0^\infty (1 - H_1(s))\,\mathrm{d}s\right)&\\ - \int_0^\infty 1 - H_1(s)\,\mathrm{d}s&\\ = (1 - H_1(t_n))\int_0^\infty 1 - H_1(s)\,\mathrm{d}s - \int_{t_n}^\infty 1 - H_1(s)\,\mathrm{d}s = O(1 - H_1(t_n))&\end{aligned}$$

Because the payoff is strictly decreasing in $\mu^{\hat{H}_1^\sigma}$, the agent would benefit from deviating to this strategy for some t_n sufficiently large, contradicting incentive compatibility. □

To prove part 2 of the Proposition 2, assume by contradiction that $H_1 \notin \mathcal{H}^{\text{NBUE}}$ and (H_0, H_1) is incentive compatible. First, by Lemma 6, either H_1 has a bounded support or (17) holds. Therefore, by Lemma 7 and Lemma 8, there exists a $t > 0$ such that

$$\left(\frac{\int_t^\infty e^{-(\rho_0+\rho_1)(s-t)} \, \mathrm{d}H_1(s)}{1 - H_1(t)}, \frac{\int_t^\infty 1 - H_1(s) \, \mathrm{d}s}{1 - H_1(t)} \right) \in \Gamma^{\text{NBUE}},$$
$$\int_0^\infty 1 - H_1(s) \, \mathrm{d}s \leq \frac{\int_t^\infty 1 - H_1(s) \, \mathrm{d}s}{1 - H_1(t)}.$$

The strategy σ that prescribes joining the queue H_1 after realizing a high lump-sum payoff and restarting (reneging and immediately rejoining H_1, arbitrarily many times) at t if he has not being served by then, by Lemma 9, induces the following statistics

$$(\delta^{\hat{H}_1^\sigma}, \mu^{\hat{H}_1^\sigma}) = \left(\frac{H_1(t)}{1 - (1 - H_1(t)) e^{-(\rho_0+\rho_1)t}} \frac{\int_0^t e^{-(\rho_0+\rho)s} \, \mathrm{d}H_1(s)}{H_1(t)}, \frac{\int_0^t 1 - H_1(s) \, \mathrm{d}s}{H_1(t)} \right),$$

and $\delta^{H_1} < \delta^{\hat{H}_1^\sigma}$, $\mu^{\hat{H}_1^\sigma} \leq \mu^{H_1}$. Because the payoff is increasing in $\delta^{\hat{H}_1^\sigma}$ and decreasing in $\mu^{\hat{H}_1^\sigma}$, the agent benefits from such a deviation, contradicting the optimality of the menu (H_0, H_1).

Part (ii)(b) is proved as a step in the proof of Lemma 12 (see Claim 4).

A.2.5 Complements to the Proof of Proposition 2

For convenience, I report a few standard results from statistics.

Definition 5. *Let X be a non-negative random variable with distribution function H and a finite mean. The mean residual life of X at t is defined as*

$$\text{MRL}(t) = \begin{cases} \mathbf{E}[X - t \mid t] & \textit{for } t < \bar{t}, \\ 0 & \textit{otherwise}, \end{cases}$$

where $\bar{t} := \sup\{t \mid 1 - H(t) > 0\}$.

Note that because X is almost surely positive, $\mathrm{MRL}(0) = \mathbf{E}[X]$, and

$$\mathrm{MRL}(t) = \frac{\int_t^\infty 1 - H(s)\,\mathrm{d}s}{1 - H(t)}.$$

Theorem 3 (Guess and Proschan, 1988, Th. 2.1)**.** *A function* MRL *is the mean residual life function of a non-degenerate at* 0 *life distribution if and only if it satisfies the following properties:*

(i) $\mathrm{MRL}\colon [0,\infty) \to [0,\infty)$.

(ii) $\mathrm{MRL}(0) > 0$.

(iii) MRL *is right-continuous (not necessarily continuous).*

(iv) $\mathrm{MRL}(t) + t$ *is increasing on* $[0,\infty)$.

(v) *when there exists* t_0 *such that* $\mathrm{MRL}(t_0^-) := \lim_{t\to t_0^-} \mathrm{MRL}(t) = 0$*, then* $\mathrm{MRL}(t) = 0$ *holds for* $t \in [t_0,\infty)$*. Otherwise, when there does not exist such a* t_0 *with* $\mathrm{MRL}(t_0^-) = 0$*, then* $\int_0^\infty 1/\mathrm{MRL}(u)\,\mathrm{d}u = \infty$ *holds.*

In addition, the mean residual life function has no downward jumps because

$$\lim_{\varepsilon\to 0^+} \mathrm{MRL}(t) - \mathrm{MRL}(t-\varepsilon) = \frac{H(t) - H(t-\varepsilon)}{1 - H(t-\varepsilon)}\mathrm{MRL}(t) - \int_{t-\varepsilon}^{t} \frac{1 - H(s)}{1 - H(t-\varepsilon)}\,\mathrm{d}s \geq 0$$

Lemma 7. *Let* X *be a nonnegative random variable with a finite mean and a bounded support that is non-degenerate at* 0*. Then, there exists* $t > 0$ *in the interior of* $\mathrm{supp}(X)$ *such that the random variable* $[X - t \mid X > t]$ *belongs to the NBUE family and* $\mathbf{E}[X - t \mid X > t] \geq \mathbf{E}[X]$.

Proof. Let $\bar{t} = \max \mathrm{supp}(X)$. Because $\mathrm{MRL}(\bar{t}) = 0$, $t^\dagger := \max\{t \mid \mathrm{MRL}(t) \geq \mu\} < \bar{t}$. By the properties of the MRL function, $\mathrm{MRL}(t^\dagger) \geq \mu$. Consequently, $[X - t^\dagger \mid X > t^\dagger]$ belongs to the NBUE family. □

Lemma 8. *Let* X *be a nonnegative random variable with a finite mean and an unbounded support that is non-degenerate at* 0*. Then, if* $\limsup_{t\to\infty} \mathrm{MRL}(t) < \mu$*, there exists a* $t > 0$ *such that the random variable* $[X - t \mid X > t]$ *belongs to the NBUE family.*

Proof. Let $t^\dagger \coloneqq \sup\{t \mid \mathrm{MRL}(t) \geq \mu\}$. By assumption $t^\dagger < \infty$ and by the properties of the MRL function, $\mathrm{MRL}(t^\dagger) \geq \mu$. Hence, $[X - t^\dagger \mid X > t^\dagger]$ belongs to the NBUE family.

□

The following result is immediate and stated without proof.

Lemma 9. *Let $\{X_j, j = 1, 2, \dots\}$ be a sequence of nonnegative independent and identically distributed random variables with finite mean and distribution function H. For $t > 0$, let*

$$Y = J \cdot t + X_J, \qquad J = \inf\{j \mid X_j < t\}.$$

Then

$$\mathbf{E}[Y] = \frac{\int_0^t 1 - H(s)\,\mathrm{d}s}{H(t)},$$

$$\mathbf{E}[e^{-(\rho_0+\rho_1)Y}] = \frac{H(t)}{1-(1-H(t))e^{-(\rho_0+\rho_1)t}}\mathbf{E}[e^{-(\rho_0+\rho_1)X} \mid X < t].$$

A.3 Proofs for Section 4

A.3.1 Preliminaries

From the proof of Lemma 2 (see Section A.2.1), the payoff from a strategy $\sigma \in \Sigma$ satisfying the properties in Proposition 1 equals (with abuse of notation) $V(\delta^{\hat{H}_0^\sigma}, \mu^{\hat{H}_0^\sigma}, \delta^{\hat{H}_1^\sigma}, \mu^{\hat{H}_1^\sigma}, \underline{p}^\sigma)$, where

$$V(\delta_0, \mu_0, \delta_1, \mu_1, p) \coloneqq \frac{\left((1-\delta_0)\frac{\rho_0}{\rho_0+\rho_1} + \delta_0 p\right)(\theta_1 - \mu_1 c) + (1-\delta_1)\frac{\rho_1}{\rho_0+\rho_1}(\theta_0 - \mu_0 c)}{\left((1-\delta_0)\frac{\rho_0}{\rho_0+\rho_1} + \delta_0 p\right)\mu_1 + (1-\delta_1)\frac{\rho_1}{\rho_0+\rho_1}\left(\mu_0 + \frac{1}{\rho_0+\rho_1}\ln\left(\frac{\rho_0}{(1-p)\rho_0 - p\rho_1}\right)\right)},$$

and the (a.s. limit of the) long-run service rate induced by σ is equal to (with abuse of notation) $\lim_{t\to\infty}\frac{1}{t}N_t = S\left(\delta^{\hat{H}_0^\sigma}, \mu^{\hat{H}_0^\sigma}, \delta^{\hat{H}_1^\sigma}, \mu^{\hat{H}_1^\sigma}, \underline{p}^\sigma\right)$, where

$$S(\delta_0, \mu_0, \delta_1, \mu_1, p) \coloneqq \frac{\delta_0 p + (1-\delta_0)\frac{\rho_0}{\rho_0+\rho_1} + (1-\delta_1)\frac{\rho_1}{\rho_0+\rho_1}}{\left(\delta_0 p + (1-\delta_0)\frac{\rho_0}{\rho_0+\rho_1}\right)\mu_1 + (1-\delta_1)\frac{\rho_1}{\rho_0+\rho_1}\left(\mu_0 + \frac{1}{\rho_0+\rho_1}\ln\left(\frac{\rho_0}{(1-p)\rho_0 - p\rho_1}\right)\right)}.$$

For later purposes, notice that the following identities hold (the decomposition in (9) is a special case of these identities):

$$
\begin{aligned}
&V(\delta_0,\mu_0,\delta_1,\mu_1,p) = \\
&\qquad S(\delta_0,\mu_0,\delta_1,\mu_1,p)\left(m(\delta_0,\delta_1,p)(\theta_1-\mu_1 c)+(1-m(\delta_0,\delta_1,p))(\theta_0-\mu_0 c)\right),
\end{aligned}
\tag{18}
$$

$$
S(\delta_0,\mu_0,\delta_1,\mu_1,p) = \frac{1}{m(\delta_0,\delta_1,p)\,\mu_1+(1-m(\delta_0,\delta_1,p))(\mu_0+t(p))}, \tag{19}
$$

where

$$
m(\delta_0,\delta_1,p) \coloneqq \frac{(1-\delta_0)\frac{\rho_0}{\rho_0+\rho_1}+\delta_0 p}{(1-\delta_0)\frac{\rho_0}{\rho_0+\rho_1}+\delta_0 p+(1-\delta_1)\frac{\rho_1}{\rho_1+\rho_0}}. \tag{20}
$$

The next lemma shows that for any tuple of summary statistics, $(\delta_0,\mu_0,\delta_1,\mu_1)$, there exists a unique optimal cutoff strategy within the class of non-reneging strategies. Let $p^*\colon \Gamma\times\Gamma\to[0,\rho_0/(\rho_0+\rho_1)]$ be defined as

$$
\begin{aligned}
&p^*(\delta_0,\mu_0,\delta_1,\mu_1) \\
&\coloneqq \begin{cases}
0 & \text{if } \beta(\delta_0,\mu_0,\delta_1,\mu_1)\le\alpha(\mu_0,\mu_1)-1<-1\,, \\
\frac{\rho_0}{\rho_0+\rho_1} & \text{if } \beta(\delta_0,\mu_0,\delta_1,\mu_1)\ge 0, \text{ or } \alpha(\mu_0,\mu_1)>0, \\
\frac{\rho_0}{\rho_0+\rho_1}\left(1-\frac{\beta(\delta_0,\mu_0,\delta_1,\mu_1)}{W_{-1}\left(e^{-1+\alpha(\mu_0,\mu_1)}\beta(\delta_0,\mu_0,\delta_1,\mu_1)\right)}\right) & \text{otherwise,}
\end{cases}
\end{aligned}
$$

where W_{-1} is the (negative branch of the) Lambert function and

$$
\begin{aligned}
\alpha(\mu_0,\mu_1) &= -(\rho_0+\rho_1)\frac{\theta_1\mu_0-\theta_0\mu_1}{\theta_1-c\mu_1}, \\
\beta(\delta_0,\mu_0,\delta_1,\mu_1) &= -\frac{\rho_0(\theta_1-c\mu_1)+(1-\delta_1)\rho_1(\theta_0-c\mu_0)}{\delta_0\rho_0(\theta_1-c\mu_1)}\,.
\end{aligned}
\tag{21}
$$

The proof of Lemma 10 is a matter of tedious algebra and omitted.

Lemma 10. *Given* $(\delta_0,\mu_0,\delta_1,\mu_1)\in\Gamma\times\Gamma$*, there exists a unique* $p\in[0,\rho_0/(\rho_0+\rho_1)]$ *that solves* $\max_{p\in[0,\rho_0/(\rho_0+\rho_1)]} V(\delta_0,\mu_0,\delta_1,\mu_1,p)$*. It equals* $p^*(\delta_0,\mu_0,\delta_1,\mu_1)$*.*

Let $V^*(\delta_0,\mu_0,\delta_1,\mu_1)\coloneqq V(\delta_0,\mu_0,\delta_1,\mu_1,p^*(\delta_0,\mu_0,\delta_1,\mu_1))$ and $S^*(\delta_0,\mu_0,\delta_1,\mu_1)\coloneqq S(\delta_0,\mu_0,\delta_1,\mu_1,p^*(\delta_0,\mu_0,\delta_1,\mu_1))$. It is easy to see that the function V^* is Gateaux differentiable in $(\delta_0,\mu_0,\delta_1,\mu_1)$ whenever $V^*(\delta_0,\mu_0,\delta_1,\mu_1)$ is strictly positive. A similar remark holds for $p^*(\delta_0,\mu_0,\delta_1,\mu_1)$ whenever interior, because it is a composition of

Gateaux differentiable functions. The following fact assembles some technical results to be used later:

Fact 1. *Let* $(\delta_0, \mu_0, \delta_1, \mu_1) \in \Gamma \times \Gamma$ *and* $p \in [0, \rho_0/(\rho_0 + \rho_1))$ *be such that* $V(\delta_0, \mu_0, \delta_1, \mu_1, p) > 0$.

(i) $V(\delta_0, \mu_0, \delta_1, \mu_1, p)$ *is strictly decreasing in* μ_0, μ_1 *and* δ_0, *and strictly increasing in* δ_1.

(ii) $p^*(\delta_0, \mu_0, \delta_1, \mu_1)$ *is increasing in* μ_0 *and* μ_1, *increasing in* δ_0 *(strictly if interior), and decreasing in* δ_1.

Lemma 10 has an immediate important consequence: the designer can always guarantee that the aggregate payoff is strictly positive by offering a single service-in-random-order queue. The result is formalized in the following lemma whose proof is omitted.

Lemma 11. *For any set of admissible parameters* $(\theta_1, \theta_0, c, \rho_0, \rho_1, \lambda) \in \mathbf{R}_{++} \times \mathbf{R}_{++} \times \mathbf{R}_{+} \times \mathbf{R}_{++} \times \mathbf{R}_{++}$, *there exists an equilibrium* $(\sigma, H) \in \Sigma \times \mathcal{H}$ *that yields a strictly positive payoff.*

A.3.2 Proof of Theorem 2

Whenever the cutoff belief in Lemma 10 is interior, it satisfies the first-order condition

$$V\left(\delta_0, \mu_0, \delta_1, \mu_1, p\right) = \frac{\theta_1 - c\mu_1}{\mu_1 + \frac{\rho_1}{\rho_0+\rho_1}\frac{1-\delta_1}{\delta_0}\frac{1}{(1-p)\rho_0 - p\rho_1}}. \tag{22}$$

For convenience, let

$$\kappa\left(\delta_0, \mu_0, \delta_1, \mu_1\right) := \frac{1-\delta_1}{\delta_0}\frac{1}{\left(1 - p^*\left(\delta_0, \mu_0, \delta_1, \mu_1\right)\right)\rho_0 - p^*\left(\delta_0, \mu_0, \delta_1, \mu_1\right)\rho_1}, \tag{23}$$

so that, when the optimal cutoff is interior, the following identity holds:

$$m\left(\delta_0, \delta_1, p^*\left(\delta_0, \mu_0, \delta_1, \mu_1\right)\right) = \frac{\frac{1}{1-\delta_1}\frac{\rho_0}{\rho_0+\rho_1} - \frac{1}{(\rho_0+\rho_1)\kappa(\delta_0,\mu_0,\delta_1,\mu_1)}}{\frac{1}{1-\delta_1}\frac{\rho_0}{\rho_0+\rho_1} + \frac{\rho_1}{\rho_1+\rho_0} - \frac{1}{(\rho_0+\rho_1)\kappa(\delta_0,\mu_0,\delta_1,\mu_1)}}. \tag{24}$$

I shall refer to these three equalities several times in the remainder of the proof, in addition to the payoff decomposition in (18) and to Fact 1.

Relaxed problem. I start by solving the relaxed program

$$\text{(RP)} \qquad \max V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right)$$

over $(\delta^{H_0}, \mu^{H_0}) \in \Gamma$ and $(\delta^{H_1}, \mu^{H_1}) \in \Gamma^{\text{NBUE}}$ subject to

$$V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right) \geq V^*\left(\delta^{H_1}, \mu^{H_1}, \delta^{H_1}, \mu^{H_1}\right), \tag{IC-θ_0}$$

$$V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right) \geq V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_0}, \mu^{H_0}\right), \tag{IC-θ_1}$$

$$S^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right) \leq \lambda. \tag{C}$$

I first state the solution of the program (RP), then conclude the proof of Theorem 2, and finally present the proof of the maximization.

Lemma 12. *There exists a solution to* (RP). *It is such that* (C) *binds and either of the following holds:*

(i) $\delta^{H_0} = \delta^{H_1}$ *and* $\mu^{H_0} = \mu^{H_1}$*;*

(ii) $\delta^{H_0} = e^{-(\rho_0+\rho_1)\mu^{H_0}}$*,* $\mu^{H_0} > \mu^{H_1}$*, and* $\delta^{H_1} = 1/(1 + (\rho_0 + \rho_1)\mu^{H_1})$*, and* (IC-$\theta_0$) *is binding.*

Conclusion of the proof of Theorem 2. It remains to prove that in the case of a separating menu, agents have no incentives to renege. Since H_0 is degenerate, restarting is suboptimal. The distribution H_1 is memoryless; hence, no agent can strictly benefit from restarting. Because $\mu^{H_1} \leq \mu^{H_1}$ and $\delta^{H_0} < \delta^{H_1}$, any agent with a belief above the invariant probability prefers H_1 to H_0; thus, even along an arbitrarily long history with no service, the high type does not find it optimal to leave his queue and join the queue H_0.

Proof of Lemma 12 First, I formulate the domain restrictions $(\delta^{H_0}, \mu^{H_0}) \in \Gamma$ and $(\delta^{H_1}, \mu^{H_1}) \in \Gamma^{\text{NBUE}}$ as explicit constraints:

$$e^{-(\rho_0+\rho_1)\mu^{H_1}} - \delta^{H_1} \leq 0, \tag{WB-1}$$

$$e^{-(\rho_0+\rho_1)\mu^{H_0}} - \delta^{H_0} \leq 0, \tag{WB-0}$$

$$\delta^{H_1} - \frac{1}{1 + (\rho_0 + \rho_1)\,\mu^{H_1}} \leq 0. \tag{EB-1}$$

Define the Lagrangian function as

$$
\begin{aligned}
L\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}, \boldsymbol{\eta}\right) = V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right) + \eta_1\left(\lambda - S^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right)\right) \\
+ \eta_2\left(V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right) - V^*\left(\delta^{H_1}, \mu^{H_1}, \delta^{H_1}, \mu^{H_1}\right)\right) \\
+ \eta_3\left(V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right) - V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_0}, \mu^{H_0}\right)\right) \\
+ \eta_4\left(\delta^{H_0} - e^{-(\rho_0+\rho_1)\mu^{H_0}}\right) + \eta_5\left(\frac{1}{1+(\rho_0+\rho_1)\,\mu^{H_1}} - \delta^{H_1}\right) \\
+ \eta_6\left(\delta^{H_1} - e^{-(\rho_0+\rho_1)\mu^{H_1}}\right),
\end{aligned}
$$

where $\boldsymbol{\eta} \in \mathbf{R}_+^6$ is a vector of multiplier. If $(\delta_0^*, \mu_0^*, \delta_1^*, \mu_1^*) \in (0,1) \times (0,\infty) \times (0,1) \times (0,\infty)$ and $\boldsymbol{\eta}^* \geq \mathbf{0}$, $\boldsymbol{\eta}^* \neq \mathbf{0}$ are such that

(i) the constraints (IC-θ_1), (IC-θ_0),(C), (WB-1), (WB-0), (EB-1) and the complementary slackness conditions are satisfied;

(ii) $L\left(\delta_*^{H_0}, \mu_*^{H_0}, \delta_*^{H_1}, \mu_*^{H_1}, \boldsymbol{\eta}^*\right) \geq L\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}, \boldsymbol{\eta}^*\right)$, for any $(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) \in (0,1) \times (0,\infty) \times (0,1) \times (0,\infty)$,

then $(\delta_*^{H_0}, \mu_*^{H_0}, \delta_*^{H_1}, \mu_*^{H_1})$ is optimal. In the following, I first derive qualitative properties of any $(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) \in (0,1) \times (0,\infty) \times (0,1) \times (0,\infty)$ satisfying (i) and (ii). It is then easy to show that for any set of parameters, such a pair exists and the optimal $(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1})$ must satisfy the conditions in the statement of Lemma 12.

First, assume, throughout the following claims, that $(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) \in (0,1) \times (0,\infty) \times (0,1) \times (0,\infty)$.

Claim 1. *If* $(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1})$ *satisfies* (IC-θ_0) *and* (IC-θ_1), $\delta^{H_0} \leq \delta^{H_1}$.

Proof. By the average cost optimality equations, there exists a unique (up to an additive constant) map $u\colon \{0,1\} \to \mathbf{R}$ and a unique $V^* \in \mathbf{R}$ such that

$$
\begin{aligned}
u(0) = \max_{(\delta,\mu)} \Bigg\{ & - V^*\mu - c\mu + \left(\delta\, p^*(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) + (1-\delta)\frac{\rho_0}{\rho_0+\rho_1}\right)(\theta_1 + u(1) - \theta_0 - u(0)) \\
& + \theta_0 + u(0) \Bigg\}, \\
u(1) = \max_{(\delta,\mu)} \Bigg\{ & -V^*\mu - c\mu + \left(\delta + (1-\delta)\frac{\rho_0}{\rho_0+\rho_1}\right)(\theta_1 + u(1) - \theta_0 - u(0)) + \theta_0 + u(0) \Bigg\},
\end{aligned}
$$

where the maxima are taken over $(\delta,\mu) \in \{(\delta^{H_0}, \mu^{H_0}), (\delta^{H_1}, \mu^{H_1})\}$. Clearly, $u(0) < u(1)$. It is easy to check that, for (IC-θ_0) and (IC-θ_1) to be satisfied, $\delta^{H_0} \leq \delta^{H_1}$. □

Claim 2. *If* $(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1})$ *satisfies* (IC-θ_0)*,* $\delta^{H_0} < \delta^{H_1}$*,* $\mu^{H_1} \leq \mu^{H_0}$*, and* $S^* (\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) < S^* (\delta^{H_1}, \mu^{H_1}, \delta^{H_1}, \mu^{H_1})$*, then* $p^*(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) = 0$.

Proof. Proceeding by contradiction, assume that $p^* (\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) > 0$. First, I show that this implies $p^* (\delta^{H_1}, \mu^{H_1}, \delta^{H_1}, \mu^{H_1}) > 0$. In fact, because the right-hand side of (22) is strictly increasing in δ^{H_0} and strictly decreasing in p, for any $h = \left(h_{\delta^{H_0}}, h_{\mu^{H_0}}, 0, 0\right)$, $h_{\delta^{H_0}} > 0$ and $h_{\mu_0} \leq 0$ such that $\nabla V^* (\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) \cdot h \leq 0$, $\nabla p^* (\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) \cdot h > 0$. As a result, (IC-$\theta_0$) implies that if $p^* (\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) > 0$, $p^* (\delta^{H_1}, \mu^{H_1}, \delta^{H_1}, \mu^{H_1}) > 0$.

Second, by (22), (IC-θ_0) implies

$$\kappa \left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right) \leq \kappa \left(\delta^{H_1}, \mu^{H_1}, \delta^{H_1}, \mu^{H_1}\right)$$
$$m \left(\delta^{H_0}, \delta^{H_1}, p^* \left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right)\right) \leq m \left(\delta^{H_1}, \delta^{H_1}, p^* \left(\delta^{H_1}, \mu^{H_1}, \delta^{H_1}, \mu^{H_1}\right)\right).$$

Since $\mu^{H_1} \leq \mu^{H_0}$, it follows that $S^* (\delta^{H_1}, \mu^{H_1}, \delta^{H_1}, \mu^{H_1}) \leq S^* (\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1})$, a contradiction. □

Claim 3. *If* $(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1})$ *solves* (RP) *and* $\mu^{H_0} < \mu^{H_1}$*, then* $p^*(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) = 0$.

Proof. First notice that $\mu^{H_0} < \mu^{H_1}$ and (IC-θ_1) imply that (IC-θ_0) is slack. Assume that $p^* (\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) > 0$, so that the first-order conditions hold. Consider a change along a direction $h = \left(0, h_{\mu^{H_0}}, 0, h_{\mu^{H_1}}\right)$, $h_{\mu_0} > 0$, $h_{\mu_1} < 0$ such that

$$m(\delta^{H_0}, \delta^{H_1}, p^*(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}))h_{\mu^{H_1}} + (1 - m(\delta^{H_0}, \delta^{H_1}, p^*(\delta^{H_0}, \mu^{H_0}, \delta_1, \mu^{H_1})))h_{\mu^{H_0}} = 0$$

If $\nabla V^* (\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) \cdot h \leq 0$, (22) implies $\nabla \kappa (\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) \cdot h > 0$ and $\nabla p (\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) \cdot h > 0$. Hence, $\nabla S^* (\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) \cdot h < 0$. As a result, one can find $h'_{\mu^{H_1}}$ such that

$$\nabla S^* \left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right) \cdot \left(0, h_{\mu^{H_0}}, 0, h'_{\mu^{H_1}}\right) = 0,$$
$$\nabla V^* \left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right) \cdot \left(0, h_{\mu^{H_0}}, 0, h'_{\mu^{H_1}}\right) > 0.$$

If instead $\nabla S^* (\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) \cdot h > 0$, let $h'_{\mu^{H_0}} > h_{\mu^{H_0}}$ be such that $\nabla S^* (\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) \cdot \left(0, h'_{\mu^{H_0}}, 0, h_{\mu^{H_1}}\right) = 0$. By the first part of the proof, $\nabla V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right) \cdot h > 0$, contradicting the optimality of $\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right)$. □

Claim 4. *If* $\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right)$ *solves* (RP), $\left(\delta^{H_0}, \mu^{H_0}\right) \neq \left(\delta^{H_1}, \mu^{H_1}\right)$, *and* $p^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right) = 0$, *then* (WB-0) *and* (EB-1) *are binding.*

Proof. Suppose first that (WB-0) is slack. In this case,

$$\begin{aligned} &\frac{\partial S^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right)}{\partial \delta_0} \\ &\qquad = -\frac{\left(\mu^{H_0} - \mu^{H_1}\right)\rho_0}{\left(1-\delta^{H_0}\right)\rho_0 + \left(1-\delta^{H_1}\right)\rho_1} \frac{\partial S^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right)}{\partial \mu_0}, \end{aligned} \tag{25}$$

and

$$\begin{aligned} &\frac{\partial V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right)}{\partial \delta_0} \\ &\qquad = -\frac{\left(\mu^{H_0}\theta_1 - \mu^{H_1}\theta_0\right)\rho_0}{\left(1-\delta^{H_0}\right)\rho_0\theta_1 + \left(1-\delta^{H_1}\right)\rho_1\theta_0} \frac{\partial V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right)}{\partial \mu_0}. \end{aligned} \tag{26}$$

Because $\theta_0 < 0 < \theta_1$,[21]

$$\frac{\left(\mu^{H_0} - \mu^{H_1}\right)\rho_0}{\left(1-\delta^{H_0}\right)\rho_0 + \left(1-\delta^{H_1}\right)\rho_1} < \frac{\left(\mu^{H_0}\theta_1 - \mu^{H_1}\theta_0\right)\rho_0}{\left(1-\delta^{H_0}\right)\rho_0\theta_1 + \left(1-\delta^{H_1}\right)\rho_1\theta_0}. \tag{27}$$

Since

$$\frac{\partial V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right)}{\partial \mu_0} < 0, \qquad \frac{\partial S^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right)}{\partial \mu_0} < 0,$$

there exists a direction $h = \left(h_{\delta^{H_0}}, h_{\mu^{H_0}}, 0, 0\right)$, $h_{\delta^{H_0}} < 0$, $h_{\mu_0} > 0$, along which

$$\nabla S^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right) \cdot h \leq 0, \qquad \nabla V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right) \cdot h > 0.$$

[21] In light of the definition of $G(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}, p)$, for any candidate optimal tuple $(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1})$, the denominator of the term on the right-hand side of equation (27) is strictly positive. For otherwise, the aggregate payoffs would be strictly negative, for any strategy that prescribes queueing with positive probability.

Because along the direction h all other constraints are either unchanged or relaxed, at the optimum, (WB-0) must bind. Assume next that (EB-1) is slack. Since,

$$\begin{aligned}&\frac{\partial S^*\left(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1}\right)}{\partial\delta_1}\\&\quad=\frac{\left(\mu^{H_0}-\mu^{H_1}\right)\rho_1}{\left(1-\delta^{H_0}\right)\rho_0+\left(1-\delta^{H_1}\right)\rho_1}\frac{\partial S^*\left(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1}\right)}{\partial\mu_1},\\&\frac{\partial V^*\left(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1}\right)}{\partial\delta_1}\\&\quad=\frac{\left(\mu^{H_0}\theta_1-\mu^{H_1}\theta_0\right)\rho_1}{\left(1-\delta^{H_0}\right)\rho_0\theta_1+\left(1-\delta^{H_1}\right)\rho_1\theta_0}\frac{\partial V^*\left(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1}\right))}{\partial\mu_1},\end{aligned}$$

by (27), there exists a direction $h=\left(0,0,h_{\delta^{H_1}},h_{\mu^{H_1}}\right)$, $h_{\delta^{H_1}}>0$ and $h_{\mu_1}>0$, along which all constraints are relaxed or unchanged and the objective function is increased. Hence, (EB-1) must bind at the optimum. □

Claim 5. *If* $(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1})$ *solves* (RP), $p^*(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1})=0$, *and* (WB-1) *is slack,* (IC-θ_0) *binds.*

Proof. Assume by contradiction that (IC-θ_0) is slack. Let $h=\left(h_{\delta^{H_0}},h_{\mu^{H_0}},0,0\right)$, $h_{\delta^{H_0}}<0$, $h_{\mu_0}>0$ be a direction such that[22]

$$h_{\delta^{H_0}}+(\rho_0+\rho_1)e^{-(\rho_0+\rho_1)\mu_0}h_{\mu_0}=0. \tag{28}$$

Assume first that $\nabla V^*(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1})\cdot h<0$. Then, by (25)–(27), $\nabla S^*(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1})\cdot h<0$. Let then $h_{\mu^{H_1}}<0$ be such that $\nabla S^*(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1})\cdot\left(h_{\delta^{H_0}},h_{\mu^{H_0}},0,h_{\mu^{H_1}}\right)=0$. Because $\partial m\left(\delta^{H_0},\delta^{H_1},p\right)/\partial\delta^{H_0}<0$, if $\nabla S^*\left(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1}\right)\cdot h=0$, $\nabla V^*\left(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1}\right)\cdot h>0$. As no constraint is violated along that direction, this contradicts the optimality of $(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1})$. Assume next that $\nabla V^*(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1})\cdot h>0$. Again, by (25)–(27), it is possible to find a direction $h'=\left(h'_{\delta^{H_0}},h'_{\mu^{H_0}},0,0\right)$ that does not violate (WB-0) and such that $\nabla S^*(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1})\cdot h=0$ and $\nabla V^*(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1})\cdot h>0$. As no constraint is violated along that direction, this contradicts the optimality of $(\delta^{H_0},\mu^{H_0},\delta^{H_1},\mu^{H_1})$. □

[22] The restriction (28) takes care of (WB-0).

Claim 2 and Claim 3 imply that at any optimal $(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1})$, $(\delta^{H_0}, \mu^{H_0}) \neq (\delta^{H_1}, \mu^{H_1})$, $p^*(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) = 0$. Hence, by Claim 4, at any optimal such a $(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1})$,(WB-0) and (EB-1) are binding. As when (EB-1) is binding, (WB-1) is slack, Claim 5 implies that (IC-θ_0) binds, which implies $\mu^{H_0} > \mu^{H_1}$. It remains to prove that (C) binds at any such optimal $(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1})$.

Claim 6. *If* $(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1})$ *is optimal,* $p^*(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) = 0$, $\delta^{H_0} < \delta^{H_1}$, *and* (IC-$\theta_0$) *binds, the constraint* (C) *binds.*

Proof. Consider a direction $h = \left(h_{\delta^{H_0}}, h_{\mu^{H_0}}, 0, 0\right)$, $h_{\delta^{H_0}} > 0$ and $h_{\mu_0} < 0$ such that $h_{\delta^{H_0}} + (\rho_0 + \rho_1)\delta^{H_0} h_{\mu_0} = 0$. It can be checked that

$$\begin{aligned}
&\frac{\partial V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right)}{\partial \delta_0} \\
&\quad = \frac{\left(\mu^{H_0}\theta_1 - \mu^{H_1}\theta_0\right)\rho_0}{\left(1-\delta^{H_0}\right)\rho_0\theta_1 + \left(1-\delta^{H_1}\right)\rho_1\theta_0} \frac{\partial V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right)}{\partial \mu_0} < 0.
\end{aligned}$$

Since $p^*(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}) = 0$,

$$\frac{\left(\mu^{H_0}\theta_1 - \mu^{H_1}\theta_0\right)\rho_0}{\left(1-\delta^{H_0}\right)\rho_0\theta_1 + \left(1-\delta^{H_1}\right)\rho_1\theta_0} < \frac{1}{(\rho_0+\rho_1)\delta^{H_0}}.$$

As a result, $\nabla V^*\left(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \mu^{H_1}\right) \cdot h > 0$. If (C) does not bind, no constraint is violated along the direction h, contradicting the optimality of $(\delta^{H_0}, \mu^{H_0}, \delta^{H_1}, \delta^{H_1})$. □

Claim 7. *If* $(\delta, \mu, \delta, \mu)$ *is optimal,* (C) *binds.*

Proof. Assume (C) does not bind. If $e^{-(\rho_0+\rho_1)\mu} < \delta$, so that (WB-1) and (WB-0) are slack, consider a direction $h = (0, h_\mu, 0, h_\mu)$, $h_\mu < 0$. The fact that $\nabla V^*(\delta, \mu, \delta, \mu) \cdot h > 0$ contradicts the optimality of $(\delta, \mu, \delta, \mu)$. Consider next the case in which $\delta = e^{-(\rho_0+\rho_1)\mu}$ and consider a direction $h = (h_\delta, 0, h_\delta, 0)$, $h_\delta > 0$. It is verified that

$$\begin{aligned}
&\frac{\partial m\left(\delta, \delta, p^*\left(\delta, \mu, \delta, \mu\right)\right)}{\partial \delta_0} + \frac{\partial m\left(\delta, \delta, p^*\left(\delta, \mu, \delta, \mu\right)\right)}{\partial \delta_1} > 0 \\
&\frac{\partial \Lambda\left(\delta, \mu, \delta, \mu, p^*\left(\delta, \mu, \delta, \mu\right)\right)}{\partial \delta_0} + \frac{\partial \Lambda\left(\delta, \mu, \delta, \mu, p^*\left(\delta, \mu, \delta, \mu\right)\right)}{\partial \delta_1} > 0.
\end{aligned}$$

Since the change $p^*\left(\delta, \mu, \delta, \mu\right)$ can be neglected (as follows from the envelope theorem), $\nabla V^*(\delta, \mu, \delta, \mu) \cdot h > 0$, yielding the desired contradiction. □

A.3.3 Proof of Theorem 1

First, I show when $c = 0$, service-in-random-order is the best discipline when the designer is constrained to offer a single non-reneging queueing discipline. Second, I show that there exist a menu that outperforms the best service-in-random-order discipline.

Let $(\delta, \mu) \in \Gamma^{\text{NBUE}}$ and $h = (h_\delta, h_\mu, h_\delta, h_\mu) \in \mathbf{R}^4_{++}$ be a direction such that $\nabla V^* (\delta, \mu, \delta, \mu) \cdot h = 0$. There are two cases. If $\rho_0\theta_1 + \rho_1\theta_0 > 0$, the optimal cutoff $p(v, \mu, \delta, \mu)$ is strictly increasing in δ and μ. Since for fixed $p, \nabla m(\delta, \delta, p) \cdot h > 0$ and $m(\delta, \delta, p)$ is strictly increasing in p

$$\nabla m (\delta, \delta, p^*(\delta, \mu, \delta, \mu)) \cdot h + \frac{\partial m (\delta, \delta, p^* (\delta, \mu, \delta, \mu))}{\partial p} (\nabla p^* (\delta, \mu, \delta, \mu) \cdot h) > 0. \tag{29}$$

By (18), $\nabla V^* (\delta, \mu, \delta, \mu) \cdot h = 0$ only if $\nabla S^* (\delta, \mu, \delta, \mu) \cdot h < 0$. Next, I show that even if $\rho_0\theta_1 + \rho_1\theta_0 \le 0$, (29) holds. To do so, I rewrite the agent's best reply problem in Lemma 10 as a choice over m instead of a choice over optimal cutoffs p. That is, given $(\delta, \mu) \in \Gamma$, there exists a unique $m \in (0, 1)$ that solves

$$\max_{m\in[\rho_0/(\rho_0+\rho_1),\rho_0/(\rho_0+(1-\delta)\rho_1))} \frac{1}{\mu + (1-m)\frac{1}{\rho_0+\rho_1} \ln\left(\frac{(1-m)\delta\rho_0}{(1-m)\rho_0-(1-\delta)m\rho_1}\right)} (m\theta_1 + (1-m)\theta_0).$$

The first-order conditions read

$$\begin{aligned} &(1-\delta)\left(m\theta_1 + (1-m)\theta_0\right)\frac{\rho_1}{\rho_0+\rho_1} \\ &-\left((\theta_1-\theta_0)\mu + \frac{\theta_1}{\rho_0+\rho_1}\ln\left(\frac{(1-m)\delta\rho_0}{(1-m)\rho_0-(1-\delta)m\rho_1}\right)\right)((1-m)\rho_0-(1-\delta)m\rho_1) = 0. \end{aligned}$$

It can be shown that the left hand side the equation above is decreasing in δ and μ, and, using the assumption $\rho_0\theta_1 + \rho_1\theta_0 \le 0$, it is increasing in m. Hence, by the implicit function theorem, along a direction $h = (h_\delta, h_\mu, h_\delta, h_\mu) \in \mathbf{R}^4_{++}$, (29) must hold, and again, by (18), $\nabla V^* (\delta, \mu, \delta, \mu) \cdot h = 0$ only if $\nabla S^* (\delta, \mu, \delta, \mu) \cdot h < 0$.

As a result, the optimal pair $(\delta, \mu) \in \Gamma^{\text{NBUE}}$ must lie at the east boundary of the set Γ^{NBUE}. Otherwise, one could increase welfare by increasing δ and μ without violating the capacity constraint.

To find the first-come first-served/service-in-random-order menu that outperforms the best single service-in-random-order queue, I show that the following system, which identifies a candidate optimal menu, has a solution:

$$\begin{aligned}
&\frac{(1-\delta_0)\frac{\rho_0}{\rho_0+\rho_1}\theta_1+(1-\delta_1)\frac{\rho_1}{\rho_0+\rho_1}\theta_0}{(1-\delta_0)\frac{\rho_0}{\rho_0+\rho_1}\frac{1-\delta_1}{(\rho_0+\rho_1)\delta_1}+(1-\delta_1)\frac{\rho_1}{\rho_0+\rho_1}\frac{\ln(1/\delta_0)}{\rho_0+\rho_1}} \\
&\qquad\qquad\qquad = \frac{\delta_1\theta_1}{\frac{1-\delta_1}{\rho_0+\rho_1}+\frac{\rho_1}{\rho_0+\rho_1}\frac{1-\delta_1}{\rho_0-(\rho_0+\rho_1)p^*\left(\frac{1-\delta_1}{(\rho_0+\rho_1)\delta_1},\delta_1,\frac{1-\delta_1}{(\rho_0+\rho_1)\delta_1},\delta_1\right)}},
\end{aligned} \tag{30}$$

$$\frac{(1-\delta_0)\frac{\rho_0}{\rho_0+\rho_1}+(1-\delta_1)\frac{\rho_1}{\rho_0+\rho_1}}{(1-\delta_0)\frac{\rho_0}{\rho_0+\rho_1}\frac{1-\delta_1}{(\rho_0+\rho_1)\delta_1}+(1-\delta_1)\frac{\rho_1}{\rho_0+\rho_1}\frac{\ln(1/\delta_0)}{\rho_0+\rho_1}}=\lambda, \tag{31}$$

$$\begin{aligned}
0 \le\ & (1-\delta_0)\,\frac{\rho_0}{\rho_0+\rho_1}\theta_1 \\
&+(1-\delta_1)\,\frac{\rho_1}{\rho_0+\rho_1}\theta_0-\rho_0\delta_0\left(\theta_1\frac{\ln(1/\delta_0)}{\rho_0+\rho_1}-\theta_0\frac{1-\delta_1}{(\rho_0+\rho_1)\delta_1}\right).
\end{aligned} \tag{32}$$

For a fixed $\delta_0 \in (0,1)$, the left-hand side of (31) is increasing in δ_1, tends to infinity as $\delta_1 \to 1$, and to 0 as $\delta_1 \to 0$. As a result, there exists a continuous curve $\mathcal{C} \subset (0,1)^2$, the first and second coordinates corresponding to δ_0 and δ_1, respectively, that is a solution to (31). It is easy to see that $\{(0,1),(1,0)\} \subset \overline{\mathcal{C}}$.

For a fixed $\delta_1 \in (0,1)$ the left-hand side of (32) is decreasing in δ_0; let $\mathcal{D}_0 \in (0,1)^2$ be the set of points that satisfy (32) as equality. Denote by D_0 the set of points lying on or above the curve $\mathcal{D}_0$. The left-hand side of (30) is increasing in δ_0 if and only if $(\delta_0,\delta_1) \in D_0$. Hence, there exists a continuous curve $\mathcal{D} \subset D_0$ that solves (30). Because $\{(0,0),(1,1)\} \subset \overline{\mathcal{D}}$, by the intermediate value theorem, the two curves $\mathcal{D}$ and $\mathcal{C}$ cross and $(\delta_0,\delta_1) \in \mathcal{C}\cap\mathcal{D} \subset D_0$ solves the system of (30)–(32).

Any solution to the system describes an incentive-compatible and feasible menu such that agents have incentive to join the first-come first-served queue as soon as their belief jumps to zero. By Lemma 13, for any $(\delta_0,\delta_1) \in \mathcal{C}\cap\mathcal{D}$,

$$m\left(\delta_0,\delta_1,p^*\left(\delta_1,\frac{1-\delta_1}{(\rho_0+\rho_1)\delta_1},\delta_1,\frac{1-\delta_1}{(\rho_0+\rho_1)\delta_1}\right)\right) < m\left(\delta_0,\delta_1,0\right),$$

which implies that

$$S^*\left(\delta_0, \frac{\ln(1/\delta_0)}{\rho_0+\rho_1}, \delta_1, \frac{1-\delta_1}{(\rho_0+\rho_1)\delta_1}\right) < S^*\left(\delta_1, \frac{1-\delta_1}{(\rho_0+\rho_1)\delta_1}, \delta_1, \frac{1-\delta_1}{(\rho_0+\rho_1)\delta_1}\right).$$

That is, the service-in-random-order discipline which yields the same payoff as a candidate menu is unfeasible. Consequently, the best feasible service-in-random-order discipline is outperformed by any candidate menu that solves (30)–(32).

Lemma 13. *Suppose* $\delta_0 \in (0,1)$ *and* $\delta_1 \in (0,1)$ *solves* (30)–(32) *and* $V^*(\delta,\ln(1/\delta_0)/(\rho_0+\rho_1), \delta_1, 1-\delta_1/((\rho_0+\rho_1)\delta_1)) > 0$. *Then,*

$$m^*\left(\delta_1, \delta_1, p^*\left(\delta_1, \frac{1-\delta_1}{(\rho_0+\rho_1)\delta_1}, \delta_1, \frac{1-\delta_1}{(\rho_0+\rho_1)\delta_1}\right)\right) < m(\delta_0, \delta_1, 0).$$

Proof. One can check that for any $\delta_1 \in (0,1)$, $p^*(\delta_1, 1-\delta_1/((\rho_0+\rho_1)\delta_1), \delta_1, 1-\delta_1/((\rho_0+\rho_1)\delta_1)) > 0$ and, by assumption, $p^*(1-\delta_1/((\rho_0+\rho_1)\delta_1), \delta_1, 1-\delta_1/((\rho_0+\rho_1)\delta_1), \delta_1) < \rho_0/(\rho_0+\rho_1)$, so the first-order condition holds. Hence,

$$\delta_0 \cdot \frac{\rho_0}{\rho_0+\rho_1} < \delta_1\left(\frac{\rho_0}{\rho_0+\rho_1} - p^*\left(\frac{1-\delta_1}{(\rho_0+\rho_1)\delta_1}, \delta_1, \frac{1-\delta_1}{(\rho_0+\rho_1)\delta_1}, \delta_1\right)\right)$$

The result follows from the definition of $m(\delta_0, \delta_1, p)$. □

A.3.4 Proof of Lemma 4

Proof of (i). First, When the candidate menu is optimal, the total payoff is bounded above by $\lambda\theta_1 - c$, which is negative for $c > \lambda\theta_1$. However, as shown after Fact 1, the designer can always guarantee that the aggregate payoff is strictly positive by offering a single service- in-random-order queue. Hence, the separating menu cannot be optimal for sufficiently high c.

Second, I show that for sufficiently high c, service-in-random-order is not optimal even if the designer is constrained to a single non-reneging queue.

Let $(\delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}})$ be the statistics of the best feasible service-in-random-order queue when the queueing cost is c. That is,

$$S^*\left(\delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}, \delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}\right) = \lambda, \qquad \mu_c^{\text{SIRO}} = \frac{1-\delta_c^{\text{SIRO}}}{(\rho_0+\rho_1)\delta_c^{\text{SIRO}}}.$$

(By Lemma 11, this system has a solution for any $c > 0$.)

Claim 8. *The following hold:*

(i) $\lim_{c\to\infty} \delta_c^{\text{SIRO}} = 1$;
(ii) $\lim_{c\to\infty} p^*\left(\delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}, \delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}\right) = \rho_0/(\rho_0+\rho_1)$.

Proof. First, $\Lambda(\delta, \mu, \delta, \mu, p)$ is strictly decreasing in p when evaluated at $p^*(\delta,\mu)$ and

$$\Lambda\left(\delta, (1-\delta)/((\rho_0+\rho_1)\delta), \delta, (1-\delta)/((\rho_0+\rho_1)\delta), p\right)$$

is strictly increasing in δ. Second, for a fixed (δ,μ), the function $p^*(\delta,\mu,\delta,\mu)$ is increasing in c and

$$p^*\left(\delta, (1-\delta)/((\rho_0+\rho_1)\delta), \delta, (1-\delta)/((\rho_0+\rho_1)\delta)\right)$$

is decreasing in δ, in both cases strictly if the cutoff belief is interior. Additionally, $p^*\left(\delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}, \delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}\right) > 0$ (see proof of Lemma 11). Hence, (i) follows by the implicit function theorem. To show (ii), notice that as $c \to \infty$, $S^*\left(\delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}, \delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}\right)$ is bounded only if $-\ln \varrho\left(p^*\left(\delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}, \delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}\right)\right) \to \infty$, i.e., if $p^*\left(\delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}, \delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}\right) \to \rho_0/(\rho_0+\rho_1)$. $\square$

For any c, let $h = (h_\delta^c, h_\mu^c, h_\delta^c, h_\mu^c) \in \mathbf{R}_+^4$ be a direction such that $\nabla V^*\left(\delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}, \delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}\right) \cdot h^c = 0$, and without loss of generality, set $h_\delta^c = 1$. I shall show that $\lim_{c\to\infty} \nabla S^*\left(\delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}, \delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}\right) \cdot h^c > 0$, which contradicts the optimality of the service-in-random-order queue. Since

$$\begin{aligned}
&\operatorname{sgn}\left(\nabla S^*\left(\delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}, \delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}\right) \cdot h^c\right)\\
&= \operatorname{sgn}\Bigg(- h_\mu^c - \frac{\ln \varrho\left(p^*\left(\delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}, \delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}\right)\right)}{\rho_0+\rho_1}\\
&\qquad\qquad \cdot \left(\frac{\partial m\left(\delta, \delta, p^*\left(\delta,\mu,\delta,\mu\right)\right)}{\partial \delta_0} + \frac{\partial m\left(\delta, \delta, p^*\left(\delta,\mu,\delta,\mu\right)\right)}{\partial \delta_1}\right)\Bigg),
\end{aligned}$$

reasoning by contradiction, if $\liminf_{c\to\infty} \nabla S^*\left(\delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}, \delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}\right) \cdot h^c \leq 0$, $\limsup_{c\to\infty} h_\mu^c = \infty$, as $-\ln \varrho\left(p^*\left(\delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}, \delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}\right)\right) \to \infty$, and the last two terms are bounded. But then, by the decomposition (19), $\nabla V^*\left(\delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}, \delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}\right) \cdot h^c < 0$, contradicting (3). It follows that $\liminf_{c\to\infty} \nabla S^*\left(\delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}, \delta_c^{\text{SIRO}}, \mu_c^{\text{SIRO}}\right) \cdot h^c > 0$.

Proof of (ii). It can be checked that $V_2^*(\delta, \mu, \delta, \mu) + V_4^*(\delta, \mu, \delta, \mu) < 0$ and $V_1^*(\delta, \mu, \delta, \mu) + V_3^*(\delta, \mu, \delta, \mu) > 0$. Hence, if the optimal pair of feasible statistics $(\delta^*, \mu^*, \delta^*, \mu^*)$ in Γ^{NBUE} lies on its west boundary, any pair of statistics $(\delta, \mu) \in \Gamma^{\text{NBUE}}$ such that $\mu < \mu^*$ (which implies $\delta^* < \delta$) must violate the feasibility constraint. As for any queueing discipline inducing no-reneging, the queue length is equal to μ/λ, the result follows.

A.3.5 Proof of Lemma 5

Let $(\delta_\lambda^{\text{FCFS}}, \mu_\lambda^{\text{FCFS}})$ be the statistics of the best feasible first-come-first-served queue when the service capacity is λ, if it exists. That is, $S^*\left(\delta_\lambda^{\text{FCFS}}, \mu_\lambda^{\text{FCFS}}, \delta_\lambda^{\text{FCFS}}, \mu_\lambda^{\text{FCFS}}\right) = \lambda$ and $\delta_\lambda^{\text{FCFS}} = e^{-(\rho_0+\rho_1)\mu_\lambda^{\text{FCFS}}}$. If such a discipline does not exists, the statement trivially holds. So, I shall focus on the case in which $(\delta_\lambda^{\text{FCFS}}, \mu_\lambda^{\text{FCFS}})$ exists along a sequence $\lambda \to \infty$. Notice that by Little's law, $\mu_\lambda^{\text{FCFS}} < 1/\lambda$. Hence, $\lim_{\lambda\to\infty} \mu_\lambda^{\text{FCFS}} = 0$. Additionally, there exists a $\overline{\mu} > 0$ such that

$$\beta\left(\mu, e^{-(\rho_0+\rho_1)\mu}\right) - \alpha\left(\mu\right) + 1 < 0,$$

and hence $p^*\left(\mu, e^{-(\rho_0+\rho_1)\mu}, \mu, e^{-(\rho_0+\rho_1)\mu}\right) = 0$ for for $0 < \mu < \overline{\mu}$. Hence, $p^*\left(\delta_\lambda^{\text{FCFS}}, \mu_\lambda^{\text{FCFS}}, \delta_\lambda^{\text{FCFS}}, \mu_\lambda^{\text{FCFS}}\right) = 0$ for $\lambda > \overline{\lambda}$. But then, for any (δ, μ) such that $p^*(\delta, \mu, \delta, \mu) = 0$ and any $h = (h_\delta, h_\mu) \in \mathbf{R}_+^2$, and in particular for $h = (h_\delta, h_\mu, h_\delta, h_\mu) \in \mathbf{R}_+^2$ such that $\nabla V^*\left(\delta, \mu, \delta, \mu\right) \cdot h > 0$, $\nabla S^*\left(\delta, \mu, \delta, \mu\right) \cdot h < 0$. Consequently, $(\delta_\lambda^{\text{FCFS}}, \mu_\lambda^{\text{FCFS}})$ is suboptimal, even if the designer is constrained to a single non-reneging queue.

References

Asmussen, Søren (2003) *Applied Probability and Queues*, New York: Springer.

Bassamboo, Achal and Ramandeep Singh Randhawa (2015) "Scheduling Homogeneous Impatient Customers," *Management Science*, 62, 2129–2147.

Bergemann, Dirk and Philipp Strack (2022) "Progressive participation," *Theoretical Economics*, 17, 1007–1039.

Bloch, Francis and David Cantala (2017) "Dynamic Assignment of Objects to Queuing Agents," *American Economic Journal: Microeconomics*, 9, 88–122.

Che, Yeon-Koo and Olivier Tercieux (2023) "Optimal Queue Design," Working Paper.

Condorelli, Daniele (2012) "What Money Can't Buy: Efficient Mechanism Design with Costly Signals," *Games and Economic Behavior*, 75, 613–624.

Cripps, Martin W. and Caroline D. Thomas (2019) "Strategic Experimentation in Queues," *Theoretical Economics*, 14, 647–708.

Garrett, Daniel F. (2016) "Intertemporal Price Discrimination: Dynamic Arrivals and Changing Values," *American Economic Review*, 106, 3275–99.

Gershkov, Alex and Benny Moldovanu (2010) "Efficient Sequential Assignment with Incomplete Information," *Games and Economic Behavior*, 68, 144–154.

Gershkov, Alex and Benny Moldovanu (2014) *Dynamic Allocation and Pricing: A Mechanism Design Approach*, Cambridge, MA: MIT Press.

Guess, Frank and Frank Proschan (1988) "Mean residual life: Theory and applications," in *Handbook of Statistics* ed. by P. R. Krishnaiah and C.R. Rao: Elsevier, Vol. 7, 215–224.

Guo, Yingni and Johannes Hörner (2020) "Dynamic Allocation without Money," TSE Working Paper 20-1133.

Hassin, Refael (1985) "On the Optimality of First Come Last Served Queues," *Econometrica*, 53, 201–202.

Jackson, Matthew O. and Hugo F. Sonnenschein (2007) "Overcoming Incentive Constraints by Linking Decisions," *Econometrica*, 75, 241–257.

Jovanovic, Boyan and Robert W. Rosenthal (1988) "Anonymous Sequential Games," *Journal of Mathematical Economics*, 17, 77–87.

Leshno, Jacob D. (2022) "Dynamic Matching in Overloaded Waiting Lists," *American Economic Review*, 112, 3876–3910.

Margaria, Chiara (2021) "Queueing to Learn," Working Paper.

Milner, Joseph M. and Tava Lennon Olsen (2008) "Service-Level Agreements in Call Centers: Perils and Prescriptions," *Management Science*, 54, 238–252.

Naor, Pinhas (1969) "The Regulation of Queue Size by Levying Tolls," *Econometrica*, 37, 15–24.

Platz, Trine Tornøe and Lars Peter Østerdal (2017) "The Curse of The First-In-First-Out Queue Discipline," *Games and Economic Behavior*, 104, 165–176.

Puterman, Martin L. (1994) *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, New York: Wiley.

Ross, Sheldon M. (1995) *Stochastic Processes*, New York: Wiley.

Scarsini, Marco and Eran Shmaya (2024) "Optimal Queueing Regimes," arXiv preprint arXiv:2401.13812.

Shaked, Moshe and J. George Shanthikumar (2007) *Stochastic orders*, New York: Springer-Verlag.

Sun, Yeneng (2006) "The Exact Law of Large Numbers via Fubini Extension and Characterization of Insurable Risks," *Journal of Economic Theory*, 126, 31–69.

Thomas, Caroline D. (2021) "Strategic Experimentation with Congestion," *American Economic Journal: Microeconomics*, 13, 1–82.