Optimal Queue Design[1]

Yeon-Koo Che[2] Olivier Tercieux[3]

December 5, 2020

Abstract

We study the optimal design of a queueing system when agents' arrival and servicing are governed by general Markov processes. The designer of the system chooses *entry* and *exit* rules for agents, their *service priority*—or *queueing discipline*—as well as their *information*, while ensuring that agents have incentives to follow the designer's recommendations not only to *join* the queue but more importantly to *stay* in the queue. The optimal mechanism has a cutoff structure—agents are induced to enter up to a certain queue length and no agents are to exit the queue once they enter the queue; the agents on the queue are served according to a first-come-first-served (FCFS) rule; and they are given no information throughout the process beyond the recommendations they receive from the designer. FCFS is also necessary for optimality in a rich domain. We identify a novel role for queueing disciplines in regulating agents' beliefs, and their dynamic incentives, revealing a hitherto unrecognized virtue of FCFS in this regard.

JEL Classification Numbers: C78, C61, D47, D83, D61
Keywords: Queueing disciplines, information design, mechanism design, dynamic matching.

1 Introduction

Consider the problem faced by somebody, called a designer, who designs a queueing system for agents seeking to receive a service or product. Agents arrive stochastically according to

[1]We are grateful to Ethan Che, Krishnamurthy Iyer, Ioannis Karatzas, Jinwoo Kim, Jacob Leshno, Vahideh Manshadi, Chris Ryan, Sara Shahanaghi, and Eduardo Teixeira, for their helpful comments. We acknowledge research assistance from Dong Woo Hahm. Yeon-Koo Che is supported by the Ministry of Education of the Republic of Korea and the National Research Foundation of Korea (NRF-2020S1A5A2A03043516)

[2]Department of Economics, Columbia University, USA. Email: yeonkooche@gmail.com.
[3]Department of Economics, Paris School of Economics, France. Email: tercieux@pse.ens.fr.

some Markov process, and are served according to another Markov process, both depending on the number of agents in the queue. Aside from these two processes, which are exogenous and thus beyond her control, the designer can choose many aspects of the queue design. She can let an arriving agent enter the queue or turn him away. She can remove an agent from the queue. She can also decide which agent will be served at each point, or more generally, how an available service capacity is allocated among agents in the queue. Finally, the designer can control how much information an agent can have about the queue or his expected waiting time throughout the process, both when he arrives at the queue and at any point after he has joined.

Casual observation of real-world queues suggests a wide range of choices available along these different dimensions of queue design. Service call centers sometimes encourage customers to wait in line (i.e., to be put on hold); other times, presumably in the face of high call volume, they tell customers to try another time. Similarly, some call centers ask customers who have already been waiting for some time, to leave the queue and return later. Queue disciplines—the service priorities for agents in the queue—also take a variety of forms: *first-come-first-served* (FCFS) is the oldest and by far the most common queue system, but *service-in-random-order* (SIRO) which assigns priority at random, has been also used. Some authors have proposed other rules such as *last-come-first-served* (LCFS) (e.g., Hassin (1985), Su and Zenios (2004), and Platz and Østerdal (2017)). Finally, a range of different information policies are commonly observed. Some queueing systems keep customers completely in the dark about the queue's length, their relative positions, or their estimated waiting times. For instance, many offices for social housing do not disclose any information on positions on waiting lists.[1] Other systems provide customers with their estimated waiting time or the number of customers ahead of them. For instance, popular ride-hailing apps provide a customer with not only the estimated arrival time of a vehicle but also its current location in a map.

A queue design, together with the primitive process, induces a Markov chain on the length of the queue; naturally we focus on the Markov chain in its steady state, or at its invariant distribution. This in turn determines an agent's expected waiting time and the service rate in steady state. For this outcome to be feasible, the associated queueing behavior must be compatible with agents' incentives. While the designer can keep an agent from joining the queue or remove one from the queue, she cannot coerce an agent to enter the queue or to stay in it against his will.[2] In other words, when recommended to either join a queue or

[1] This is the case, for instance, for several housing choice voucher programs in California, e.g., PCCDS Housing Service or HACA among others.

[2] Given our assumption, a possible incentive by an agent to queue up against desires of the designer is not a problem, since the designer can simply block an entry by such an agent. In this sense, excessive entry, as observed by Naor (1969) to occur under FCFS, does not arise our model. However, this feature of our model is natural and realistic in many "regulated" settings where a designer has control over entry into a queue.

stay in the queue, an agent must have an incentive to obey this recommendation given the information he has. The designer chooses a feasible outcome that maximizes a weighted sum of agents' welfare and the service provider's profit, the former consisting of a lump-sum benefit from service minus linear waiting costs, and the latter consisting of a fixed (possibly shadow) profit it enjoys for each agent served. Since the weight is arbitrary, the designer can be a service provider maximizing profit, a consumer advocate maximizing agents' welfare, or a regulator who values both.

The question is: *how should the designer choose all different aspects of the queue design?* Our answer is strikingly simple and consistent with many observed practices of queue design. The optimal queue design has the following features:

(i) The optimal queue design has a *cutoff* policy: namely, there exists a maximal queue length $K \geq 0$ such that agents are recommended to enter the queue as long as its length is less than $K-1$ but they are turned away if the queue reaches K in length. No agents leave the queue once they join it.[3]

(ii) Those who join the queue are then prioritized to receive a service according to FCFS.

(iii) No information is provided to agents beyond the recommendations they receive to join or to stay in the queue.[4]

Result (i) means that no agent should be either removed or incentivized to leave the queue once he or she joins it, except possibly when the queue is full. In other words, reneging—or abandonment of the queue—is never part of the optimal queue behavior, again except possibly when the queue is full. Results (ii) and (iii) mean that, at least in the canonical model we consider, the most tried-and-true queueing norm is as good as any others and often strictly better, provided that agents are kept in the dark throughout.

The intuition behind the information policy—no information beyond recommendation—is explained as follows. It is well known and intuitive that incentive constraints are relaxed most when agents are given as little information as possible. If an agent has the incentive to join a queue or to stay in it for a set of signals he may receive under any queueing discipline, he must also have the incentive to join the queue when all these signals are pooled into one. Since this "pooled" signal is precisely what the agent will have given "no information"

[3]When the queue length is $K-1$, an agent is recommended to enter with positive probability possibly equal to one. If this probability is less than one, the entry is "rationed" at $K-1$. While such an rationed entry can be equivalently implemented by an agent exiting at a positive rate when the queue is full, we "normalize" the policy throughout so that no agents exit.

[4]Since recommendations contain information about the state, this policy should not be confused with "no information" authors often use, which refers to "no communication" what so ever. Agents can make Bayesian inferences on their expected waiting times, based on the recommendation they receive, the queue design that the designer commits to, and the elapsed time after joining the queue.

beyond the recommendation to enter a queue or to stay in it, the same queue discipline must also be incentive compatible under a no information policy. This explains why providing no information is optimal.

To explain why FCFS is optimal, fix an optimal entry and exit policy—i.e., a cutoff policy with some maximal length K. Assuming agents obey the recommendation, this induces a distribution of queue length in the steady state. Since our agents are homogeneous, the expected waiting time *when averaged across possible initial queue lengths* is the same for each agent, and does not depend on the queueing discipline in use. Then, given no information, the incentive for joining the queue will be the same across all queueing disciplines, meaning if it is incentive compatible to join the queue under any queue discipline it will be incentive compatible to do so under any other. Hence, from this standpoint, FCFS is not particularly necessary or desirable.

However, the dynamic incentives that agents face—particularly their incentive to obey the recommendation to stay in the queue once they enter it—differ across queueing disciplines, assuming the no information policy. The reason is that the belief an agent forms about the remaining waiting time evolves differently over time under different queueing disciplines. Our main insight is that, under a very mild condition on the primitive process, the evolution of this belief becomes progressively more favorable under FCFS. Consequently, under this mild condition, agents are willing to stay in the queue under FCFS with no information. Thus, the optimal queueing outcome can be implemented by FCFS, provided that no information is provided to the agents.

The progressively favorable beliefs under FCFS stem from its fundamental property: namely, that one's service priority can only improve over time under FCFS. Hence, starting with any initial queue length, the elapse of time is indeed *good news* about the remaining waiting time: as time passes, one's priority, and correspondingly, one's chance of receiving service by a given date, can only improve under FCFS. But this is not the only force at work. Since an agent is not told about the queue length k when he joins the queue (recall *no information*), his belief about this will be also updated as time progresses. On this account, the elapse of time is actually *bad news*, since it indicates that the agent likely underestimated the initial length of the queue when he joined it. Our proof boils down to showing that the good news dominates the bad news. As noted above, this means that incentive compatibility is maintained throughout once an agent joins under FCFS.

The same is not true, however, for other queueing disciplines, as the belief evolution need not be as favorable over time. Consider SIRO. Since priority is assigned randomly, the more agents there are in the queue, the less likely it is for an agent to receive service. When no information is given, an agent's belief about the queue length is therefore crucial for his dynamic incentives. This is what the agent updates over time to decide whether it is worth staying in the queue. Since service priority does not improve with time under SIRO, the updating is not as favorable as under FCFS. Further, elapse of time without being served indicates that there are more agents in the queue than he initially thought. So the agent is more likely to become pessimistic as time passes. Indeed, we can find simple cases such as the standard $M/M/1$ queue in which the belief worsens over time to such a degree that an agent leaves the queue after entering it, thus failing the incentive requirement necessary for implementing the optimal cutoff policy.

In sum, we identify a novel role played by queueing disciplines in shaping and regulating agents' beliefs about the remaining waiting time, and thus their dynamic incentives to stay in a queue—a crucial element for the eventual outcome when agents cannot be coerced to remain in a queue. In particular, our analysis reveals a hitherto unrecognized virtue of FCFS in this regard. Further, we show that no other queueing disciplines possess this virtue. Specifically, for *any* queueing discipline differing from FCFS, one can always find an environment under which it is strictly suboptimal no matter the information policy adopted. In this sense, FCFS is not only optimal under the no information policy, it is also *necessary* for optimality in a rich domain.

The current paper follows the long line of queueing theory research, in particular, the *rational queueing* literature. This literature, which has developed into a significant body of work since the seminal work by Naor (1969), studies the strategic behavior of rational Bayesian agents in a variety of queueing scenarios.[5] While sharing their focus and approach, the current paper is distinguished from standard works in this literature in several respects.

First, our Markovian model is general and flexible enough to encompass many settings of interest. A typical queueing model tends to focus on a specific process such as $M/M/1$ or $M/M/c$. Similarly, a standard dynamic matching model in economics considers a specific match technology. By contrast, our model allows the arrival and servicing of agents to follow general Markov processes that may depend on the current queue length, which nests $M/M/c$ (which in turn subsumes $M/M/1$) queueing models as well as recent economic models of dynamic matching as special cases.[6]

Second, we consider agents' incentives not only to *join* but more importantly to *stay* in the queue when recommended by the designer to do so. Addressing these latter dynamic incentives distinguishes the current paper from most of existing ones. There are a few papers that consider incentives by agents to abandon a queue, or to "renege"; see Hassin and Haviv (1995), Haviv and Ritov (2001), Mandelbaum and Shimkin (2000), Sherzer and Kerner (2018), and Cripps and Thomas (2019). However, these papers approach the issue as a positive theory, trying to explain reneging as a rational strategic response to various

[5]See Hassin and Haviv (2012) and Hassin (2012), for an excellent survey of the literature.

[6]As we mention in Section 3, the models in this literature consider agents who can only be matched with agents (or objects) in the queue that are compatible with them. Assuming that each pair is compatible with some fixed probability, the effective arrival rate (i.e., the rate at which an agent joins the queue) and the effective service rate (i.e., the rate at which an agent leaves the queue) depends on the number of agents in the queue.

features such as nonlinear waiting costs or aggregate uncertainty. Our approach is instead to treat the issue from a normative perspective, and more systematically as part of incentive design, following the tradition of mechanism design, as we note next.

Third, the current paper is distinguished in its comprehensive treatment of many aspects of queueing system design. Most of the existing papers do not consider the optimal entry/exit policies explicitly, but rather focus on some, typically unregulated, exogenously given queueing environment. Likewise, queueing disciplines are often assumed in the literature to be FCFS or, less frequently, SIRO. In addition, agents' information is also typically fixed; authors often assume that agents have either full information about the queue or no information whatsoever.[7]

A few papers study the optimal design of queueing disciplines while taking other aspects of queueing system as given. Following Naor (1969)'s seminal observation that FCFS causes agents to queue *excessively*, ignoring the "congestion" externality they inflict on later agents, Hassin (1985) and Su and Zenios (2004) argue that LCFS can "cure" this externality and is optimal for agents.[8] Meanwhile, Leshno (2019) considers a different environment in which FCFS creates *too few* incentives for queueing (due to excess supply of agents), and finds that other mechanisms such as SIRO would perform better by providing greater incentives for queueing. The suboptimality of FCFS in these papers rests crucially on the unregulated nature of queueing system and the full information assumption (i.e., complete observability of queues), respectively. We consider a general model wherein agents' private incentives for queueing may either exceed or fall short of the designer's objective and show that FCFS is always optimal if she can regulate excessive queue and choose the optimal information policy—namely, *no information*. Further, its optimality is strict if one also considers agents' dynamic incentive to stay in the queue after joining it, which the above papers do not consider.[9]

[7] Here no information means that agents truly do not have *any* information, including a recommendation from the designer. In fact, this assumption is typically made in the context of an *unregulated* environment, where there is no designer or supervisory entity overseeing or managing the queue. See Hassin and Haviv (2012) for the canonical description of the unregulated environment.

[8] Platz and Østerdal (2017) find a similar result when there are a continuum of agents who enter at their endogenously chosen times. See also Haviv and Oz (2016) for alternative schemes in the observable environment and Haviv and Oz (2018) for extensions to the unobservable queue environment.

[9] Several papers study alternative queueing disciplines in environments that are less related or comparable to ours. FCFS is shown to be optimal in Bloch and Cantala (2017) and a part of the optimal design in Margaria (2020) in models where, unlike the standard queueing model, the lengths of queues are non-stochastic, either because arrival occurs only when an agent exits (the former) or because there are a continuum of agents (the latter). Further, they do not consider information design, so the reason for the optimality of FCFS is totally different in these models than in our model. Kittsteiner and Moldovanu (2005) consider the allocation of priority in queues via bidding mechanisms where processing time is private information. The crucial difference is the use of transfers implicit in bidding mechanisms, which is not allowed in our model. Also related is the growing literature on dynamic matching; see Akbarpour, Li, and Gharan

Despite its practical relevance, information design has received attention only recently in the queueing literature; see Simhon, Hayel, Starobinski, and Zhu (2016), Lingenbrink and Iyer (2019), and Anunrojwong, Iyer, and Manshadi (2020).[10] While the latter two papers identify the same optimal information design as the current paper, they do not study the optimal queueing discipline but they instead take FCFS as given. These models are also more special in their primitive processes, and they do not consider dynamic incentives. By contrast, we allow all these dimensions of design—entry, exit, queueing disciplines, and information design—to be chosen optimally by the designer (in the face of the dynamic incentive problem).

In terms of style, the current paper is closest in spirit to the mechanism design literature, pioneered by Myerson (1981), which takes the underlying physical character of the environment as given but otherwise allows the designer to choose all other aspects of the system optimally.[11] Interestingly, our main findings are also similar in flavor as those of Myerson (1981): *the optimal mechanism is both simple and resembles commonly observed practices.* As mentioned, the cutoff policy conforms to the standard practice of capping the queue length at some level. The optimality of FCFS accords well with its prevalent use in practice, and is reassuring in light of its perceived fairness (see Larson (1987)). The *no information beyond recommendation* policy also conforms to standard practice in call centers which put customers on hold, often with no information on their waiting times, unless they are *explicitly* discouraged from waiting on line.[12]

The rest of the paper is organized as follows. Section 2 describes the model and formulates the optimization problem facing the designer. Section 3 discusses the scope of applications and relationships with some existing models. Section 4 solves the relaxed design problem and

(2020), Akbarpour, Combe, Hiller, Shimer, and Tercieux (2020), Baccara, Lee, and Yariv (2020), Doval and Szentes (2018), and Ashlagi, Burq, Jaillet, and Manshadi (2019), among others. The primary focus of this literature is the optimal timing of matching and assignment, rather than queueing incentives. None of these papers consider information design or the incentives to stay in the queue, but they have elements (e.g., preference heterogeneity) not present in our model. Hence, the current work should be viewed as complementary to this literature.

[10] In a less related model, Ashlagi, Faidra, and Nikzad (2020) study optimal dynamic matching with information design, showing that FCFS, together with an information disclosure scheme, can be used to implement the optimal outcome. Although similar at first glance, their model is quite different from, and not easily comparable to, ours. There are a continuum of agents in their model, and their information policy pertains to the quality of good rather than to agents' queue position. In particular, the virtue of FCFS in regulating agents' beliefs on where they stand in the queue is orthogonal to Ashlagi, Faidra, and Nikzad (2020)'s insights.

[11] Within the queueing literature, the optimal design or control literature focuses on the ex ante choice of the service and arrival process, which takes as given (see Shaler Stidham (2009) for a survey). Hence, one can view the current work as complementing this literature.

[12] As we already pointed out, offices for social housing often provide applicants with very limited information on their position in the list. In addition, these offices often close waiting lists when they are too long.

establishes the cutoff structure of the optimal mechanism. Section 5 shows that FCFS with no information can attain the solution of the relaxed problem and is thus optimal. Section 6 show how other information policies or other queueing rules (even under no information) such as SIRO are suboptimal. Section 7 show that FCFS with no information is the only design that is optimal for all regular environments. Section 8 concludes.

2 Model and Preliminaries

We consider a general queueing model in which agents arrive sequentially at a queue to receive a service. Time is continuous, with $t \in \mathbb{R}_+$.

Agents' payoffs. There are three parties: a *designer*, who organizes resource allocation including the queueing policy, a *service provider* who services agents, and *agents* who receive service. As will be seen, the designer may be the service provider, a representative of the agents, or a planner who reflects the welfare of both parties.

The agents are homogeneous in their preferences. Each agent enjoys a payoff of

$$U(t) \triangleq V - C \cdot t,$$

if she receives service after waiting $t \in \mathbb{R}_+$ periods, where $V > 0$ is the net surplus from service (possibly after paying a fixed service fee to the designer) and $C > 0$ is a per-period cost of waiting. The service provider earns revenue $R > 0$ from each agent she services.[13] The designer's objective, which will be specified more fully below, is a weighted sum of the service provider's and agents' payoffs. An agent's outside option, which she collects when not joining the queue or from exiting one, yields a payoff normalized to zero.

Primitive process. At each instant, given the number of agents in the queue, or **queue length**, $k \in \mathbb{Z}_+$, an agent arrives at a Poisson rate of $\lambda_k > 0$ and an agent in a queue (if it is nonempty) is served at a Poisson rate of $\mu_k > 0$. Hence, a pair (λ, μ), where $\lambda \triangleq \{\lambda_k\}$ and $\mu \triangleq \{\mu_k\}$, specifies a **primitive process**. We assume that μ_k is uniformly bounded over k. The state-contingent primitive process allows for a general class of environments that nests a variety of settings of interest, including $M/M/c$ queue models as well as dynamic matching models, which are of growing interest within economics (see Section 3 for further details on the scope of applications).

Designer's policy. The designer has at her disposal a number of instruments. We focus on an anonymous stationary Markovian policy that treats all agents identically based on a

[13] In many contexts, R may not take a monetary form but rather a shadow benefit of fulfilling a service requirement (e.g., promised by an outsourced call center).

couple of **state** variables: the queue length k and the queue position ℓ, namely the arrival order of an agent among those in a queue. The stationarity restriction means that the policy does not depend on the calendar time.

The designer chooses a set of policies:

- **entry rule** $x = \{x_k\}$, where $x_k \in [0,1]$ is the probability with which an arriving agent is asked to join a queue of length k. We let $\mathcal{X}$ denote the set of all entry rules.

- **exit rule** $y = \{y_{k,\ell}\}$, where $y_{k,\ell} \geq 0$ is the Poisson rate at which an agent with queue position ℓ is removed from a queue of length k. The exit rule y captures both the explicit policy of diverting some agent away from a service pool, as has been considered by Mandelbaum and Shimkin (2000), as well as the abandonment induced by a queueing policy (to be described below). We let $\mathcal{Y}$ denote the set of all exit rules.

- **queueing rule** $q = \{q_{k,\ell}\}$, where $q_{k,\ell} \geq 0$ is the Poisson rate at which an agent receives service when the queue length is k and her position in the queue is ℓ. Service capacity is limited by the service rate μ_k, so any queueing rule must satisfy $\sum_{\ell=1}^{k} q_{k,\ell} = \mu_k$, for all k. We let $\mathcal{Q}$ denote the set of all such feasible queueing rules. The set $\mathcal{Q}$ is large enough to encompass all standard queueing disciplines, such as **first-come-first-served (FCFS)**, in which case $q_{k,\ell} \triangleq q^F_\ell$ for some q^F_ℓ which depends only on ℓ; **last-come-first-served (LCFS)**, in which case $q_{k,\ell} \triangleq q^L_{k-\ell}$ for some $q^L_{k-\ell}$ which depends only on $k-\ell$; **service-in-random-order (SIRO)**, in which case $q_{k,\ell} \triangleq q^S_k$ for some q^S_k which depends only on k. In fact, we can allow queueing rules to be fully general, i.e., without limiting ourselves to those that depend only on (k,ℓ); examples include rules that allow service probabilities to vary with time and to depend on other variables. However, our class entails no loss since the optimal rule in the general case belongs to this class.

- **information rule** $I = \{I_t\}_{t \in \mathbb{R}_+}$, where I_t represents the information an agent has about the state—i.e., the queue length k and her position ℓ—after staying in the queue for a duration $t \geq 0$. Note that an information rule pertains to an agent who has just arrived and is deciding whether to join the queue (in which case $t = 0$) as well as to those who are already in the queue ($t > 0$). As is well-known, say from Kamenica and Gentzkow (2011), the information I_t can be represented as the distribution of "posterior beliefs" about (k,ℓ) when recommended to join the queue (i.e., for $t = 0$) and/or to stay in the queue (for $t > 0$), i.e., $I_t \in \Delta(\Delta(\mathbb{Z}_+ \times \mathbb{Z}_+))$. The belief is a mean-preserving spread of the true state, adapted to the filtration generated by the primitive process (λ,μ) and a Markov policy (x,y,q). We let $\mathcal{I}$ denote the set of all information rules. The set $\mathcal{I}$ is general enough to include all realistic information structures that are feasible. Special cases are **full information**, in which case I_t coincides with the distribution of the true state (k,ℓ), and **no information** (beyond recommendations), in which case the conditional belief I_t is degenerate on the posterior obtained by the Bayes updating via (x, y, q), from the prior beliefs I_0. As with the queueing rule, given the constructive nature of our result, we can allow the class of information structures to be fully general, beyond those that refer to the beliefs on (k, ℓ) over time.[14] We limit the class of information structures just for expositional ease.

Given the primitive process (λ, μ), a Markov policy (x, y, q) generates a Markov chain—more specifically, a birth-and-death process—on the queue length k. Given (λ, μ), we only consider Markov policies that induce a Markov chain with an invariant distribution. As is well-known, the invariant distribution $p \triangleq (p_0, p_1, ...p_\infty)$ is characterized by the following balance equation:

$$\lambda_k x_k p_k = (\mu_{k+1} + \sum_{\ell} y_{k+1,\ell}) p_{k+1}, \; \forall k \in \text{supp}(p) \tag{B}$$

The LHS of the equation is the rate with which the queue length transits from k to $k+1$: with probability p_k the queue length is k, in which case an agent arrives at rate λ_k and is recommended to join the queue with probability x_k. The balance equation (B) requires this rate to equal the rate with which the queue length transits from $k+1$ to k, namely its RHS: with probability p_{k+1} the queue length is $k+1$, in which case an agent is served at rate μ_{k+1} or is removed with rate $\sum_\ell y_{k+1,\ell}$ from the queue. We say that an entry/exit policy $(x, y) \in \mathcal{X} \times \mathcal{Y}$ **implements** an invariant distribution p if (x, y, p) satisfies (B), and call such a triple (x, y, p) an **outcome**.

Incentives. We assume that the designer may prevent an agent from joining the queue and may also remove an agent from a queue, but that the designer cannot coerce an agent to join the queue or to stay in a queue against her will. In this sense, the entry rule (specifying if agents join the queue) and exit rule (specifying if agents stay in the queue) should be regarded as a *recommendation*. In other words, when recommended to enter the queue or to stay in the queue, an agent must have an incentive to obey that recommendation, given the information available to her. For conceptual clarity, we describe this obedience constraint indirectly through the *expected waiting time* induced by the policy. Suppose that our policy (x, y, q) induces a state-contingent expected waiting time $\tau_{k,\ell} \in \mathbb{R}_+$, namely, *the expected number of periods an agent in state (k, ℓ) must wait before she either receives the service or is removed from the queue.* Note that, given the linearity of preference in the waiting time, only the expected waiting time is payoff relevant; all other aspects of the waiting time distribution do not matter. The policy also induces the probability $\{Y_{k,\ell}\}$ of eventually being removed from the queue before receiving service. Finally, the policy (x, y, q) together with

[14] For instance, the class of relevant information structures may include those that refer to beliefs about the particular service probabilities at each public or even private history.

information rule I induces posterior beliefs $\gamma^t = \{\gamma^t_{k,\ell}\}$ for an agent who is recommended to join or to stay in the queue.

The obedience constraints require that any agent recommended to join the queue or to stay in the queue must have an incentive to follow that recommendation:

$$\sum_{k,\ell} \gamma^t_{k,\ell} \left[V(1 - Y_{k,\ell}) - C \cdot \tau_{k,\ell}\right] \geq 0, \forall \gamma^t \in \text{supp}(I_t), \forall t \geq 0. \tag{IC}$$

In the sequel, we refer to the incentive constraint for t by (IC_t). We say that a queueing/information policy $(q, I) \in \mathcal{Q} \times \mathcal{I}$ **implements** an outcome (x, y, p) if (IC) holds, meaning the induced beliefs distributions and the exptected waiting times satisfy (IC_t) for all $t \geq 0$. Even though we interpret an implemented outcome as resulting from the designer's policy choice, this is without loss, due to the revelation principle. Our model can capture any equilibrium outcome, both regulated and unregulated. For instance, consider the textbook unregulated and unobservable $M/M/1$ queue governed by FCFS, in which agents make their entry decisions without any recommendation or any information about the queue length, which is unobserved (see Hassin and Haviv (2012) for instance). It is an equilibrium for an agent to enter the queue with some probability $e \in (0, 1]$ and stay in the queue forever once joining the queue.[15] In our model, this corresponds to our entry policy of $x_{k,\ell} = e$ and $y_{k,\ell} = 0$, for all k, ℓ (along with FCFS and no information).

Problem statement. The designer's objective is evaluated at the stationary distribution $p = (p_k)$ of the Markov chain. It can be written as follows:

$$W(p) \triangleq (1 - \alpha)R \sum_{k=1}^{\infty} p_k \mu_k + \alpha \sum_{k=1}^{\infty} p_k(\mu_k V - kC),$$

where $\alpha \in [0, 1]$. The first term is the flow expected profit for the service provider: with probability p_k, the queue has k agents, and an agent is served at rate μ_k in that case, generating a "profit" of R for a fulfilled service. In a customer service context, the profit may not take the form of monetary fees collected from customers but rather the shadow value of fulfilling a warranty service or more generally addressing any customer needs. The second term is the flow expected utility for agents: again with probability p_k, the queue has k agents, each of whom pays flow holding/waiting cost of C (the second term), and an agent is served at rate μ_k, in which case he receives a surplus of V (gross of the waiting cost). The objective is a weighted sum of these two terms, with weight $\alpha \in [0, 1]$.

[15] More specifically, if λ is sufficiently low (more precisely, if $(\mu - \lambda)V > C$), then agents enter with probability $e = 1$ (since, by Lemma S2 in the online appendix, (IR) holds when the queue length is infinite if and only if $(\mu - \lambda)V > C$). If not, then there exists a random entry probability $e \in (0, 1)$ such that if all agents adopt this mixing strategy, each agent becomes indifferent to entry, making it an equilibrium behavior.

The designer's problem is to choose $(p, x, y, q, I) \in \Delta(\mathbb{Z}_+) \times \mathcal{X} \times \mathcal{Y} \times \mathcal{Q} \times \mathcal{I}$ to

$$[P] \qquad \text{Maximize } W(p)$$
$$\text{subject to (B) and } (IC),$$

where the expected waiting times $\{\tau_{k,\ell}\}$ and exit rates $\{Y_{k,\ell}\}$ in (IC) are induced by (p, x, y, q). In words, the designer picks the outcome that maximizes her objective among those that are implementable by some queueing/information policy. While the entry/exit policy (x, y) uniquely pins down the invariant distribution, we include p as part of a choice. Let $\mathcal{W}$ denote the supremum of the value of program $[P]$.

3 Scope of Applications

Our model encompasses a variety of queueing and dynamic matching models considered by the existing literature.

• *$M/M/1$ queue model:* This is the most canonical queueing model in which the arrival rate λ_k and service rate μ_k do not depend on the queue length k. Naor (1969) and Hassin (1985) investigate agents' incentives to join the queue under FCFS and LCFS, respectively. Hassin and Haviv (1995) consider "reneging," or agents' dynamic incentives to leave the queue, when they face nonlinear holding costs in an unobservable and unregulated system operated by FCFS. More recently, Simhon, Hayel, Starobinski, and Zhu (2016) and Lingenbrink and Iyer (2019) study information design to manage agents' incentive to join the queue, both under FCFS queueing rule. While similar to these latter papers in considering information design, our model is more general in several respects: the queueing environment (we allow for state-dependent arrival and service rates), agents' incentives (we consider their incentives to stay in, not just to join, a queue), and queueing rules (we allow for fully general queueing rules, not just FCFS).

• *$M/M/c$ queue model:* This generalizes the $M/M/1$ queue model to allow for multiple $c \geq 2$ servers, each with exponential service time. As with M/M/1, the arrival rate is independent of the queue length k, but the service rate is linear in the number of available servers, so $\mu_k = \min\{k, c\}\mu$. Haviv and Ritov (2001) extend Hassin and Haviv (1995)'s inquiry about reneging incentives to the $M/M/c$ setup.

• *Team servicing model:* Suppose there are m customers (or machines) each having a servicing need arising at an independent Poisson rate while operating (see Gnedenko and Kovalenko (1989), p. 42). There are c servers which can serve a customer at rate μ. When there are k agents in the queue, the new arrival rate is $\lambda_k = (m - k)\lambda$ and the service occurs at rate $\mu_k = \min\{k, c\}\mu$.

• *Dynamic one-sided matching with stochastic compatibility:* Suppose each agent is compatible with another agent with probability $\theta \in (0, 1]$. In this model, an agent joins a queue only when he arrives at some rate η *and* is incompatible with the agents already in the queue (or else he matches and leaves the queue), which occurs with probability $(1-\theta)^k$, and an agent leaves the queue when he matches, which occurs with probability $\eta(1-(1-\theta)^k)$. This is a special case of our model in which $\lambda_k = \eta(1-\theta)^k$ and $\mu_k = \eta(1-(1-\theta)^k)$. Doval and Szentes (2018) consider such a model with $p = 1$ and study agents' incentive to join a queue under FCFS. Akbarpour, Li, and Gharan (2020) study the limit as $\theta \in (0, 1)$ tends to 0 but the arrival rate increases. Their focus differs from ours; for instance, they do not consider the incentive to join or stay in a queue, the queueing rule, or information design. Instead, they study the benefit from thickening the market, which we do not consider.

• *Dynamic two-sided matching with stochastic compatibility:* Heterogeneous agents on one side match with heterogeneous agents or objects (e.g., housing) on the other side. If the types of the matched pair are compatible, then high surplus is realized; if not, a low surplus is realized. The designer operates buffer queues for different types of agents or objects to keep the agents waiting until a compatible match is found. Leshno (2019) and Baccara, Lee, and Yariv (2020) consider such models. In these models, if one buffer queue is active, the other is empty. Hence, the system can be analyzed as a one-dimensional Markov chain, and so it can be embedded in our setting. Some of our results below rely on the system induced by a given policy to exhibit birth and death processes. Indeed, this feature is satisfied under the optimal policy under Baccara, Lee, and Yariv (2020) but not under Leshno (2019). Nevertheless, our central results apply to the latter setup. Baccara, Lee, and Yariv (2020) consider optimal matching policy under both FCFS and LCFS, whereas Leshno (2019) considers a general class of queueing rules, and finds random queueing rules to be optimal. Both papers assume complete information, i.e., neither considers information design. Again, the current paper is differentiated by its consideration of broad incentive issues (i.e., the incentive to stay in, not just to join, a queue) and a general class of queueing rules as well as information design. The fact that we draw a different conclusion on the optimal queueing rule—namely, FCFS—relative to Leshno (2019) is attributed to our ability to combine information design with the choice of a queueing rule (see Section 8 for further discussion).

4 Optimality of the Cutoff Policy

The designer's problem $[P]$ is in general difficult to solve. Instead, we are interested in the relaxed problem facing the designer. Consider a linear program:

$$[P'] \qquad \max_{p \in \Delta(\mathbb{Z}_+)} W(p)$$

subject to

$$\sum_{k=1}^{\infty} p_k(\mu_k V - kC) \geq 0; \tag{IR}$$

$$\lambda_k p_k - \mu_{k+1} p_{k+1} \geq 0. \tag{B$'$}$$

Here, the planner maximizes the designer's objective subject only to the individual rationality (IR) and a weakening (B′) of the balance equation (B).

It is clear that $[P']$ is a relaxation of $[P]$. First, (IR) must be implied by (IC). If the former condition fails in a mechanism, the agents do not collectively break even. Then, there must exist *some* agent and *some* belief induced by that mechanism such that the agent with that belief would not wish to join a queue when called upon to do so. Hence, (IC) would fail.[16] Next, since the $(x_k, y_{k,\ell})$ are nonnegative, (B) implies (B′). Let $\mathcal{W}^*$ denote the supremum of the value of program $[P']$. Then, whenever $\mathcal{W}^* < \infty$, we must have $\mathcal{W}^* \geq \mathcal{W}$.

The program $[P']$ is interesting in its own right: it can be interpreted as the problem facing a planner who chooses the invariant distribution p directly to maximize her objective, simply facing the primitive process (λ, μ), but disregarding agents' incentives altogether, except for guaranteeing some minimal payoff for them. Ultimately, however, we are interested in $[P']$ as an analytical tool for characterizing an optimal queue design that solves $[P]$, since a solution to this relaxed program $[P']$ may be attained by a mix of policy tools (x, y, q, I).

Indeed, our ultimate goal is to prove such a policy mix exists, which will then imply that it optimally solves $[P]$, the real objective of interest. The analysis will involve demonstrating three claims: (i) an optimal solution p^* to $[P']$ exists, (ii) there exists an entry/exit policy (x^*, y^*) that implements p^*; namely, the associated outcome (x^*, y^*, p^*) satisfies (B); and (iii) there exists a queueing/information policy (q^*, I^*) that implements the optimal outcome (x^*, y^*, p^*), meaning $(x^*, y^*, p^*, q^*, I^*)$ satisfies (IC). The remainder of this section will address (i) and (ii), while claim (iii) will be taken up in the next section.

With respect to (ii), we establish not only that the optimal p^* is well-defined and can be implemented by some entry/exit policy (x^*, y^*), but also that, under a very mild condition

[16] This can be shown more precisely. Fix any (x, y, p, q, I) that satisfies (IC_0). Aggregating (IC_0) across all beliefs $\gamma^0 \in \text{supp}(I_0)$, we get

$$\int_{\gamma^0 \in \text{supp}(I_0)} \sum_{k,\ell} \gamma^0_{k,\ell}[V(1 - Y_{k,\ell}) - C\tau_{k,\ell}] I_0(d\gamma^0) \geq 0.$$

Clearly, the *ex ante* probability of receiving service, $\int_{\gamma^0 \in \text{supp}(I_0)} \sum_{k,\ell} \gamma^0_{k,\ell}(1 - Y_{k,\ell}) I_0(d\gamma^0)$, equals $\sum_k p_k \mu_k / [\sum_k p_k \lambda_k x_k]$—the average rate of receiving service divided by the average rate of entering the queue at p. Next, by Little's law, the *ex ante* expected waiting time, $\int_{\gamma^0 \in \text{supp}(I_0)} \sum_{k,\ell} \gamma^0_{k,\ell} \tau_{k,\ell} I_0(d\gamma^0)$, equals $\sum_k p_k k / [\sum_k p_k \lambda_k x_k]$—the average queue length divided by the average entry rate. Subsituting these two expressions and simplifying the terms, the above inequality implies (IR).

on (λ, μ), (x^*, y^*) takes a particularly intuitive form:

Definition 1. An entry/exit policy (x, y) is a **cutoff policy** if there exists $K \in \mathbb{Z}_+ \cup \{+\infty\}$ such that $x_k = 1$ for all $k = 0, 1, ..., K-2$, $x_{K-1} \in (0, 1]$, and $x_k = 0$ for all $k \geq K$ and that $y_{k,\ell} = 0$ for all k, ℓ.

In words, under a cutoff policy p, the designer sets a maximum queue length K and recommends an arriving agent to join a queue as long as $k \leq K-1$ and recommends those who join the queue to stay in the queue (i.e., never to exit) until they are served. In other words, no agent is diverted away or induced to abandon his queue, once he has joined it. It is possible that $x_{K-1} \in (0, 1)$ in which case the K-th entrant may be randomly rationed.[17] Although a cutoff policy seems natural, it may not arise without policy intervention. For instance, the aforementioned unregulated/unobservable $M/M/1$ queue (governed by FCFS) will not exhibit a cutoff structure if agents randomize on entry with an interior probability $e \in (0, 1)$.[18]

Observe that an invariant distribution p is implemented by a cutoff policy (x, y) with maximal length K (potentially infinite) if and only if $\text{supp}(p) = \{0, ..., K\}$ and (B′) binds for all $k = 0, ..., K-2$ and holds for $k = K-1$ (with weak inequality). It is this latter feature that we will establish below. To this end, we define a condition.

Definition 2. The service process $\mu = \{\mu_k\}$ is **regular** if (i) μ_k is nondecreasing in k and (ii) $\mu_k - \mu_{k-1}$ is nonincreasing in k.

Regularity of the service process is very mild. In fact, condition (i) follows from a simple optimality axiom: the maximal service rate μ_k facing k agents must be no less than the maximal service rate μ_{k-1} facing $k-1$ agents, since the service provider can simply ignore one agent. Condition (ii) is less obvious but also very reasonable; in fact, without (ii), FCFS is not well defined, as will be seen later (see Lemma 1). For instance, in an $M/M/1$ queue, both λ_k and μ_k are constant, trivially satisfying both conditions. In an $M/M/c$ queue, λ_k is constant in k and $\mu_k = \min\{k, c\}\mu$ satisfies (i) and (ii). Likewise, with the random compatibility model, $\lambda_k = \eta(1-\theta)^k$ and $\mu_k = \eta(1-(1-\theta)^k)$, and so (i) and (ii) are satisfied.

Our characterization follows:

Theorem 1. An optimal solution of $[P']$ exists. If μ is regular, there is an optimal solution of $[P']$ implemented by a cutoff policy with maximal queue length $K^* \geq \arg\max_k \mu_k V - kC$.

[17] While we assume $y_k = 0$ for all k, this is just a convenient normalization. If $x_{K-1} \in (0, 1)$ in a cutoff policy, the same p^* can be implemented by any (x', y') such that $x'_{K-1} = \frac{\mu_K + \sum_\ell y'_{K,\ell}}{\mu_K} x_{K-1}$; see (B). In this sense, the reader should interpret the cutoff policy as an equivalence class involving a set of such pairs. This means that while it is unnecessary to induce an agent to exit from a queue after he joins it, doing so when the queue is full is consistent with a cutoff policy if $x_{K-1} \in (0, 1)$. In other words, encouraging a customer to come back later is not at odds with a cutoff policy.

[18] Recall from Footnote 15 that such a mixing is an equilibrium if $(\mu - \lambda)V < C$.

Proof. See Appendix A. ∎

The intuition behind the result is most clear when the coefficient on p_k in the objective as well as in (IR) is decreasing in $k \geq 1$—this occurs, for instance, if μ_k is constant in k as in the $M/M/1$ queue model. In this case, one can increase the value of the objective and relax (IR) by shifting probability mass p_k toward lower values of k in the sense of first-order stochastic dominance while keeping constant the mass at state 0.[19,20] The simple intuition is that adding an agent when the queue is long entails more wait (for agents collectively) than when the queue is short. This logic suggests that the mixing equilibrium in the aforementioned unregulated/unobserved $M/M/1$ queue is suboptimal for any α.[21] The agents will be collectively better off if agents are encouraged to enter fully if $k < K^*$ but never if $k > K^*$, for some K^*. In short, a cutoff policy is optimal.

To show this for a general "regular" service process, we use the fact that the coefficient $f(k) \triangleq \mu_k((1-\alpha)R + (\alpha+\xi)V) - (\alpha+\xi)ck$ of p_k in the Lagrangian function, for any Lagrangian multiplier $\xi \geq 0$ for (IR), is single-peaked when μ_k is regular: namely, f is increasing when $k < k^*$, is constant when $k^* \leq k \leq k^{**}$, and is decreasing when $k > k^{**}$, where k^* and k^{**} are possibly zero or infinite, in which case f is monotonic. The single peakedness of f means that on the increasing region, one would like to put the greatest possible mass on a higher k within that region—a goal that is accomplished by binding (B′) for k in that region. For k in the decreasing region, the intuition provided above applies.

The proof requires some care since the Lagrangian method may not be valid in an infinite dimensional LP. Our approach follows several steps. First, we consider a finite K-dimensional LP—one in which $p_k = 0$ for all $k > K$—and prove by using the Lagrangian method that its optimal solution $p^K = (p_0^K, ..., p_K^K)$ exhibits a cutoff policy. Second, we show that an optimal solution $\overline{p} = (\overline{p}_0, ..., \overline{p}_\infty)$ to $[P']$ exists. This follows from the observation that the set of p's satisfying (IR) is closed and its elements have a vanishing tail sum so that the feasible set of solutions is "tight" and therefore sequentially compact, by Prokhorov's theorem. This, together with the upper semi-continuity of the objective, gives us the existence of an optimal solution. Third, the same observation means that the value of K-truncation of $\overline{p}$, $\overline{p}^K \triangleq$

[19] Note that this does not mean that reducing the queue length all the way down to zero is desirable, for the value of objective is zero at $k = 0$. Maintaining a nonzero queue length is desirable from the pure efficiency standpoint. If the queue length is too low, there is a risk of server(s) going idle and wasted (which happens when the agents in the queue are served in a row without a new arrival). However, the decreasing value of objective means that the benefit from reducing this risk decreases as more agents are added to the queue.

[20] The only remaining issue then is that the LP, which is infinite dimensional, admits an optimal solution. The space of p's satisfying (IR) and with a cutoff structure is "tight" and is thus sequentially compact by the Prokhorov's theorem. Lingenbrink and Iyer (2019) uses this method to arrive at a similar conclusion. However, this simple approach does not work for our theorem due to the greater generality.

[21] As is seen from Lemma S2 in the online appendix, the mixing occurs in equilibrium if $(\mu - \lambda)V < C$ since in that case (IR) is violated if the queue length is infinite.

$(\overline{p}_0, ..., \overline{p}_K)$, which is feasible for $[P']$, converges to $\mathcal{W}^*$ as $K \to \infty$. Fourth, by definition, p^K attains a weakly higher value than $\overline{p}^K$, so its limit p^* as $K \to \infty$ attains $\mathcal{W}^*$. Finally, the set of feasible p exhibiting a cutoff policy is closed, so the limit p^*, which is optimal, retains the cutoff structure.

It is possible that the optimal cutoff policy may have $K^* = \infty$. In that case, it is optimal for agents to enter into the queue irrespective of its size and no agent is ever induced to abandon the queue; an unregulated queue would then implement this optimal outcome (if it is incentive compatible—an issue we address in the next section). However, we next show that this possibility is somewhat limited.

Proposition 1. Assume μ is regular. If $\alpha > 0$, then the optimal cutoff policy has a queue length $K^* < \infty$. If $\alpha = 0$, the optimal cutoff policy has $K^* < \infty$ if and only if $\sum_{k=1}^{\infty} \prod_{\ell=1}^{k} \frac{\lambda_{\ell-1}}{\mu_\ell} [\mu_k V - Ck] < 0$. If the latter condition holds, then (IR) binds if $\alpha > 0$ is sufficiently small or if $\alpha < 1$ and R is sufficiently large. Finally, if (IR) binds, then generically $x^*_{K^*-1} \in (0, 1)$.[22]

Proof. See Appendix B. ∎

The optimality of a finite queue length is not surprising when $\alpha > 0$: an arbitrarily long queue would entail an arbitrarily large marginal loss on the designer's objective, so reducing the length will improve the objective without violating (IR). If $\alpha = 0$, so the designer maximizes the service provider's profit, then $K^* = \infty$ is possible as long as (IR) is not violated. The latter requires the arrival rate to be sufficiently small relative to the service rate and for C to be sufficiently small relative to V. For instance, in the $M/M/1$ model, $K^* = \infty$ if and only if $(\mu - \lambda)V \geq C$ (see Corollary S1 in the online appendix).

The last result implies that, while randomizing on the entry or exit of an agent is never optimal for $k < K^* - 1$, it is generically optimal for the threshold state $K^* - 1$ when (IR) is binding. While this may be surprising at first glance, it follows from the binding (IR) and the use of a cutoff policy.[23]

5 Optimality of FCFS with No Information

In this section, we establish the general optimality of FCFS with minimal information. Assume that the service rate is regular. Then, by Theorem 1, the optimal solution p^* to $[P']$ is implemented by a cutoff policy (x^*, y^*) with a maximal queue length $K^* \in \mathbb{Z}_+ \cup \{+\infty\}$.

[22] The expression "generically" means that $\sum_{k=1}^{K} \prod_{\ell=1}^{k} \frac{\lambda_{\ell-1}}{\mu_\ell} [\mu_k V - Ck] \neq 0$, for all $K \in \mathbb{Z}_+ \cup \{\infty\}$.

[23] The use of a cutoff policy means (B′) is binding for all $k = 0, ..., K^*-2$, which pins down $p^*_k = 1, ..., K^*-1$, once p^*_0 is determined. Hence, the fact that (IR) holds with equality together with the fact that the p^*_k's sum to one requires randomization at $k = K^* - 1$, except in a knife-edge case. The arguments in Appendix B provide a precise characterization for the rationing rule together with $p^* = (p^*_k)$.

Recall that the optimal cutoff policy may involve random entry at $k = K^* - 1$; we let $x^*_{K^*-1} \in (0, 1]$ denote the optimal randomization at $k = K^* - 1$. To avoid the trivial case, we assume that $K^* > 1$.

In the sequel, we fix the optimal outcome (x^*, y^*, p^*) and $K^* > 1$ the maximal queue length. We will show that FCFS, together with an optimal information design, implements (x^*, y^*, p^*); namely, (IC) holds under that policy. Since $[P']$ is a relaxation of $[P]$, this will prove that the identified policy mix solves $[P]$.

First, we formally define FCFS as follows:

Definition 3. A queueing rule is first-come-first-served (FCFS) if, for each $\ell = 1, ..., K$, there exists q^*_ℓ such that $q_{k,\ell} = q^*_\ell$ for all $k \geq \ell$.

In other words, under the FCFS queueing rule, the probability that an agent is served/matched depends only on his queue position, namely the number of agents who arrived before that agent and are still in the queue. Recall the feasibility requirement $\sum_{\ell=1}^{k} q_{k,\ell} = \mu_k$ for all k, which now simplifies to $\sum_{\ell=1}^{k} q^*_\ell = \mu_k$, for each $k = 1, ..., K^*$. This pins down the service probability:

$$q^*_\ell = \mu_\ell - \mu_{\ell-1}, \forall \ell = 1, ..., K^*,$$

where $\mu_0 \equiv 0$. Part (ii) of the regularity condition then implies that $q^*_i \geq q^*_j$ if $i \leq j$, which captures the intuitive property that an agent with a low-numbered position, i.e., one who arrived earlier, has a higher chance of being served than an agent with a higher-numbered position, i.e., one who arrived later. It is useful to consider a few examples. In the $M/M/1$ queue model, $\mu_k = \mu$ for all $k \geq 1$, for some $\mu > 0$. In this case, $q^*_1 = \mu$ and $q^*_\ell = 0$ for all $\ell > 1$, so only the agent at the top position can be served at any time. In the $M/M/c$ queue model, there are $c > 1$ servers, so $\mu_k = \min\{c, k\}\mu$ for all $k \geq 1$, for some $\mu > 0$. In this case, $q^*_\ell = \mu$ for all $\ell \leq c$, and $q^*_\ell = 0$ for all $\ell > c$, so agents with queue positions no greater than the number of servers can be served at any time. Under the dynamic matching model with stochastic compatibility, we have $q^*_\ell = \eta(1-\theta)^{\ell-1}\theta$, which is clearly decreasing in ℓ.

Naturally, under FCFS, the expected waiting time depends only on one's queue position ℓ, so we use τ^*_ℓ to denote the expected waiting time for an agent with queue position ℓ. Given the primitives, this can be pinned down exactly.

Lemma 1. For any $\ell = 1, ..., K^*$, $\tau^*_\ell = \ell/\mu_\ell$. Given the regularity of μ, τ^*_ℓ is nondecreasing in ℓ. If $2\mu_1 > \mu_2$, then τ^*_ℓ is strictly increasing in ℓ.

Proof. See Appendix C.1 ∎

We next introduce the information rule I^*, wherein no information is provided to each agent both at the time of joining the queue and after joining the queue, beyond what he can infer from the recommendations to join or stay in the queue. This means that when he joins at time $t = 0$, he holds his prior belief on his position ℓ based on the underlying invariant distribution. From then on, he updates the belief on his queue position at each $t > 0$ according to Bayes rule without any further information.

Given the queueing and information rules (q^*, I^*), the incentive constraint at time t is given by

$$(IC_t) \qquad V - C\sum_{\ell=1}^{K^*} \tilde{\gamma}_\ell^t \cdot \tau_\ell^* \geq 0,$$

where $\tilde{\gamma}^t = (\tilde{\gamma}_1^t, ..., \tilde{\gamma}_{K^*}^t) \in \Delta(\{1, ..., K^*\})$ is the belief on his position in the queue after spending time t on the queue. Since the expected waiting time depends only on one's position, the belief on other variables such as the queue length k does not affect the agent's incentive to join or stay in the queue.

Given the information rule I^*, the belief at the time of joining the queue must be:

$$\tilde{\gamma}_\ell^0 = \begin{cases} \frac{p_{\ell-1}\tilde{\lambda}_{\ell-1}}{\sum_{i=0}^{K^*-1} p_i \tilde{\lambda}_i} & \text{if } \ell = 1, ..., K^* \\ 0 & \text{if } \ell > K^*, \end{cases} \tag{1}$$

where $\tilde{\lambda}_k$ is an "effective" arrival rate given by: $\tilde{\lambda}_k \triangleq \lambda_k$ for $k = 0, ..., K^* - 2$, and $\tilde{\lambda}_{K^*-1} \triangleq x_{K^*-1}^* \lambda_{K^*-1}$. This corresponds to the probability of being at position ℓ conditional on entering the queue. Its numerator is the probability that an agent joins the queue in state $\ell - 1$, which is the probability of the state being $\ell - 1$ (i.e., that there are $\ell - 1$ agents already in the queue) multiplied by the probability of entry per unit time in that state $\tilde{\lambda}_{\ell-1}$. Its denominator is the total probability of entering the queue per unit time.

The next lemma observes that the candidate policy (q^*, I^*) provides the agents with incentives enter the queue:

Lemma 2. The queueing/information policy (q^*, I^*) satisfies (IC_0).

Proof. See Appendix C.2. ∎

In order to ensure (IC_t) holds under (q^*, I^*) for each $t > 0$, we need to understand how an agent's belief evolves once he joins the queue. Since the optimal queue requires that no agent abandon the queue, (IC_t) requires that agents' beliefs about their queue positions should become (at least weakly) more favorable as time passes.

Suppose that an agent has belief $\tilde{\gamma}^t$ after spending time $t \geq 0$ in the queue. At time $t + dt$, by Bayes rule, his belief is updated to:

$$\tilde{\gamma}_\ell^{t+dt} = \frac{\tilde{\gamma}_\ell^t (1 - \sum_{i=1}^{\ell} q_i^* dt) + \tilde{\gamma}_{\ell+1}^t \sum_{i=1}^{\ell} q_i^* dt}{\sum_{i=1}^{K^*} \tilde{\gamma}_i^t (1 - q_i^* dt)} + o(dt).$$

The numerator is the probability that his queue position is ℓ after staying in the queue for length $t + dt$ of time. This event occurs if either (i) the agent already has position ℓ in the queue at time t and none of them, including himself, are served during time increment dt; or (ii) if he has position $\ell + 1$ at t and one agent ahead of him is served by $t + dt$. The denominator in turn gives the probability that the agent has not been served by time t. Hence, given that an agent has not been served by t, the above expression gives the conditional belief that his position in the queue is ℓ at time $t + dt$.

We can use the feasibility requirement $\sum_{i=0}^{\ell} q_i^* = \mu_\ell$ to rewrite the belief updating rule as follows:

$$\tilde{\gamma}_\ell^{t+dt} = \frac{(1 - \mu_\ell dt)\tilde{\gamma}_\ell^t + \mu_\ell dt \tilde{\gamma}_{\ell+1}^t}{\sum_{i=1}^{K^*} \tilde{\gamma}_i^t (1 - q_i^* dt)} + o(dt). \tag{2}$$

We now study how the belief updates dynamically over time under (q^*, I^*). The statistic we focus on is the **likelihood ratio** $r_\ell^t \triangleq \frac{\tilde{\gamma}_\ell^t}{\tilde{\gamma}_{\ell-1}^t}$ of beliefs about being in queue position ℓ versus being in queue position $\ell - 1$ after spending time t on the queue. One can use (2) to derive a system of ordinary differential equations (ODEs) on the likelihood ratios:

$$\dot{r}_\ell^t = r_\ell^t \left(\mu_{\ell-1} - \mu_\ell - \mu_{\ell-1} r_\ell^t + \mu_\ell r_{\ell+1}^t\right),$$

where $\ell = 2, ..., K^*$. Further, the invariant distribution p^* can be used to obtain the boundary conditions, $r_\ell^0 = \frac{\tilde{\lambda}_{\ell-1}}{\mu_{\ell-1}}$, for $\ell = 2, ..., K^*$, where $\tilde{\lambda}_k$ is an effective arrival rate: $\tilde{\lambda}_k \triangleq \lambda_k$ for $k < K^* - 1$ and $\tilde{\lambda}_{K^*-1} \triangleq x_{K^*-1}^* \lambda_{K^*-1}$. Appendix C.3 derives this system of ODEs and establishes existence of a unique solution.

We provide a condition under which these likelihood ratios—the solution to the above ODEs—decline over time, meaning one's belief about his position becomes progressively favorable under (q^*, I^*). At first glance, this seems obvious under FCFS since, starting at any position ℓ at time $t = 0$, an agent's queue position can *only* improve as time passes. Yet, this is not the only event about which the agent updates his beliefs. The agent is also updating his belief about his initial position ℓ. The elapse of time without being served is "bad" news in this regard, as it suggests that he may have been too optimistic about his position initially, causing him to revise his initial queue position pessimistically.

Indeed, the bottom curve in the graph of Figure 1 (which assumes an $M/M/1$ queue with $K^* = 2$) shows that one's belief that his initial queue position was $\ell = 1$ falls over time.[24] The difference between the top and bottom curves in the graph of Figure 1 is explained by

[24] More precisely, the bottom graph plots

$$\Pr\{\ell = 1 \text{ at } t = 0 | \text{ not served by } t\} = \frac{\tilde{\gamma}_1^0 e^{-\mu t}}{\tilde{\gamma}_1^0 e^{-\mu t} + \tilde{\gamma}_2^0 (2e^{-\mu t} - e^{-2\mu t})},$$

where $\tilde{\gamma}_1^0 e^{-\mu t}$ is the joint probability that the queue was empty when he joined it and he is not served by time t, and $\tilde{\gamma}_2^0 (2e^{-\mu t} - e^{-2\mu t})$ is the probability that the queue had one agent when he joined it and he is not served by t.

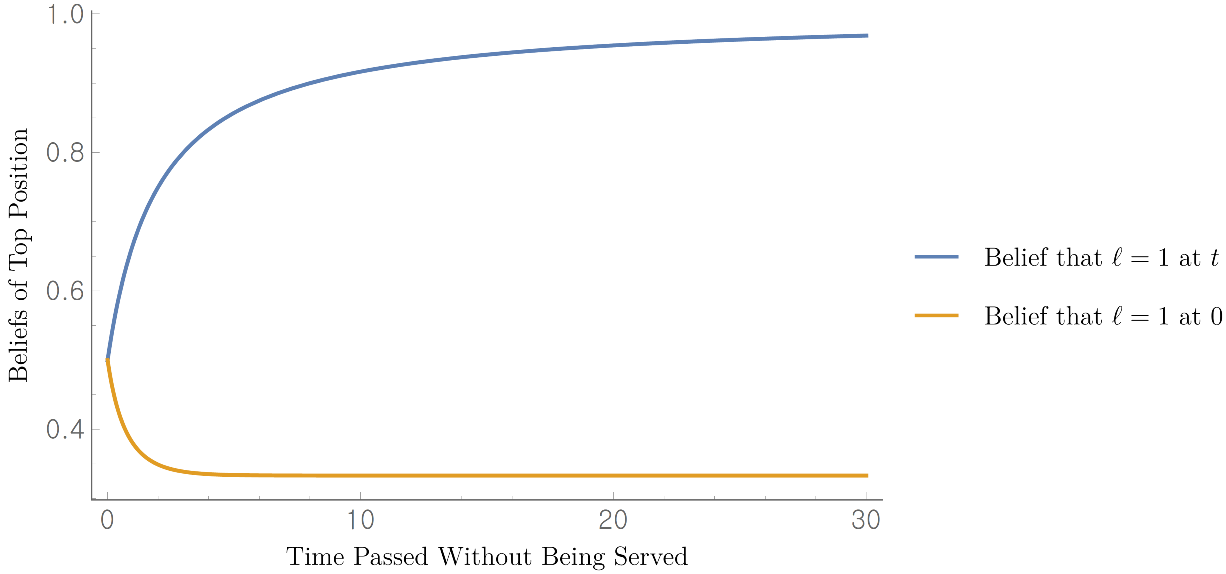

Figure 1: Belief about position $\ell = 1$
Note: $M/M/1$ with $K^* = 2$; $\lambda = \mu = 1$.

a growing belief that an initial position of $\ell = 2$ has moved up to $\ell = 1$. One can see that this latter "position-improvement" effect dominates the worsening posterior about the initial position, resulting in an overall improvement in the belief, shown by the top curve.

In fact, this result holds generally under a mild condition:

Definition 4. The primitive process (λ, μ) is **regular** if the service process μ is regular and if $\lambda_k - \lambda_{k-1} \leq \mu_k - \mu_{k-1}$ for each $k \geq 1$.

In all examples of Section 3, λ_k is either constant or decreasing, so the regularity of μ implies that (λ, μ) is regular. The next lemma provides the key result:

Lemma 3. Assume that the primitive process (λ, μ) is regular. Then, for all $\ell \in \{2, ..., K^*\}$, r_ℓ^t is nonincreasing in t for all $t \geq 0$.

Proof. See Appendix C.3. ∎

Hassin and Haviv (1995) and Haviv and Ritov (2001) establish a similar result; they show that, given FCFS, an agent's waiting time exhibits an increasing failure rate (so his waiting is increasingly likely to stop) over time under the unregulated $M/M/1$ and $M/M/c$ queue models, respectively. This result arises primarily from agents' employing a strategy of queueing only for a finite time in these models, which is in turn a rational response to the nonlinear waiting cost (i.e., "a deadline" effect).[25] No such nonlinear waiting costs are

[25] Without the nonlinear waiting costs, one can show that in an unregulated/unosbservable $M/M/1$ queue

assumed in our model. Instead, our current result arises under the optimal cutoff strategy, and under a general birth-death process, not just $M/M/1$ or $M/M/c$. Our proof method also differs from the standard argument, which focuses on establishing an increasing hazard rate of the service commencement. Our argument instead focuses on how agents' beliefs evolve over time, given a general birth-and-death process. This approach has some advantages. It reveals competing forces in the belief updating. For instance, if, contrary to regularity, $\lambda_k - \lambda_{k-1} > \mu_k - \mu_{k-1}$ for all k, then a delay is more of a signal about the initial queue length being long than about predecessors have been served, so the pessimistic updating about the initial position may dominate the optimistic updating about a shrinking queue. In that case, agents may renege immediately after joining the queue, causing (IC_t) to fail for a small $t > 0$.

We are now in a position to prove the following theorem.

Theorem 2. Assume that the primitive process is regular. Then, FCFS with no information (q^*, I^*) implements the optimal outcome (x^*, y^*, p^*). Consequently, $(x^*, y^*, p^*, q^*, I^*)$ is an optimal solution of $[P]$.

Proof. It suffices to prove that (IC_t) holds for all $t \geq 0$. Note first that, by Lemma 2, (IC_0) holds. Next consider (IC_t) for any $t > 0$. Lemma 3 proves that $r_\ell^t \leq r_\ell^0$ for each ℓ. Since τ_ℓ^* is nondecreasing in ℓ by Lemma 1, this means that

$$\sum_{\ell=1}^{K^*} \tilde{\gamma}_\ell^t \cdot \tau_\ell^* \leq \sum_{\ell=1}^{K^*} \tilde{\gamma}_\ell^0 \cdot \tau_\ell^*,$$

so we have

$$V - C\sum_{\ell=1}^{K^*} \tilde{\gamma}_\ell^t \cdot \tau_\ell^* \geq V - C\sum_{\ell=1}^{K^*} \tilde{\gamma}_\ell^0 \cdot \tau_\ell^* \geq 0,$$

where the last inequality follows from (IC_0). Hence, (IC_t) holds for any $t > 0$. ∎

We make two remarks. First, the above result relies on the designer's ability to stop an agent from entering a queue. While the designer does have such a power in many settings, the power is unnecessary if $V\mu_{K^*} \leq CK^*$, which holds for instance if (IR) is binding at the optimal outcome. In that case, the designer can simply issue a "recommendation" not to enter (when $k = K^*$), and the agent will obey that recommendation.[26] Second, FCFS with no information does not just implement the optimal outcome, but it can also implement any outcome (x, y, p) satisfying (IR) and (B) in which x_k is nonincreasing in k and $y_{k,\ell} = 0$ for

(that the former paper considers), beliefs about one's queue position stays constant over time. This means that the "good" news one gets from the elapse of time about position improvement precisely cancels out the "bad" news he gets about the initial queue position!

[26] Given the length K^* (which the agent infers from the recommendation not to enter), he expects to wait for $\tau_{K^*}^* = K^*/\mu_{K^*}$ (recall Lemma 1).

each ℓ, k. Essentially, such an outcome guarantees the regularity of the "effective" processes $(\tilde{\lambda}, \mu)$, where $\tilde{\lambda}_k \triangleq x_k \lambda_k$, which allows Theorem 2 to apply.[27]

6 Other Common Queue Designs

Theorem 2 shows that under a mild condition, namely the regularity of the primitive process, FCFS with *no information* is optimal. This leaves open the possibility that other queue designs may also be optimal in some situations. To see this, suppose for instance $\alpha = 1$ so that the designer maximizes the welfare of the agents. Then, we can show that Naor (1969)'s classic result (obtained for the $M/M/1$ queue) generalizes to our more general Markov process: *agents would have excess incentives to queue under FCFS with full information.*[28] Since our designer can simply stop excessive queueing, this means that FCFS, and similarly SIRO, with full information, denoted I^{FI}, can be used to attain optimum. Of course, *no information*, which is optimal, would have also worked in that case. In fact, one can show that, with *no information* I^*, agents will always have the incentives under the optimal cutoff policy *regardless of the queue discipline.*

Proposition 2. Suppose first $\alpha = 1$ and μ is regular. Then, FCFS or SIRO with full information, (q^*, I^{FI}), implements the optimal cutoff outcome (x^*, y^*, p^*). Next, fix any $\alpha \in [0, 1]$ and the optimal cutoff outcome (x^*, y^*, p^*). Then, the *no information* policy I^* satisfies (IC_0) under any feasible queueing rule q.

Proof. See Appendix D.1. ∎

The second result establishes the *equivalence* of alternative queueing rules under *no information*, in terms of agents' incentives to enter a queue (but not necessarily in terms of their incentive to stay in the queue). The intuition behind this result is simple. Given *no information*, agents' incentives to join a queue depends only on the expected waiting time they anticipate when they join. By Little's law, the expected waiting time does not depend on the queueing rule and, by (IR), the waiting time is short enough for entry to be profitable. This result suggests that, given *no information*, alternative queueing rules could matter only because of the dynamic incentives they generate, namely, their incentives to *stay* after they have joined the queue.

Indeed, despite this positive result for some queue designs, we will show more generally that commonly-used queue designs that differ from our optimal design are often suboptimal in many situations. We first show that "no information" is indispensable for optimality when (IR) binds.

[27] Appendix C.3 provides the argument to extend Lemma 3.

[28] This generalization, which could be of independent interest, is shown in Appendix D.1.

Proposition 3. Suppose next (IR) binds at (x^*, y^*, p^*) and the information rule $I \in \mathcal{I}$ induces a non-degenerate distribution of beliefs at $t = 0$ leading to distinct expected waiting times. Then, for any $q \in Q$, (q, I) fails (IC) at $t = 0$ and cannot implement the optimal policy.

Proof. See Appendix D.2. ∎

We now study specific queue designs in detail. In the sequel, we will focus on situations in which incentive problems are most demanding, namely when (IR) is binding at the optimal cutoff policy that solves $[P']$. As noted in Proposition 1, this is the case when α is sufficiently small or when $\alpha < 1$ and R is sufficiently large. In that case, the optimal cutoff policy involves a finite cap $K^* < \infty$. From now on, we fix such an optimal outcome (x^*, y^*, p^*).

6.1 FCFS with Information.

We know from Lemma 1 that the expected waiting time is k/μ_k for an agent if he joins a queue of length $k - 1$ under FCFS. The hypothesis of the second part of Proposition 3—the one leading to the suboptimality of providing information—will be met if $1/\mu_1 < K^*/\mu_{K^*}$ and the queue length is fully observable or if $1/\mu_1 < 2/\mu_2$ and the information rule admits signals that are ordered in *first-order stochastic dominance* (FOSD).[29] Examples of the latter information design include scenarios in which customers are told estimated waiting times or some information about the number of customers ahead of them.

Corollary 1 (FCFS with Information)**.** If $1/\mu_1 < K^*/\mu_{K^*}$ then FCFS with full information cannot implement the cutoff policy. If $1/\mu_1 < 2/\mu_2$ and the information rule I admits $\gamma, \gamma' \in \text{supp}(I_0)$, where $\gamma' \neq \gamma$ are ordered in FOSD, then FCFS with I cannot implement the optimal policy.

Proof. If $1/\mu_1 < K^*/\mu_{K^*}$, then, under full information, the expected waiting time when the initial queue has nobody ahead of the agent and the expected waiting when the initial queue has $K^* - 1$ agents ahead of him are distinct, so the result follows from Proposition 3. Next, if $1/\mu_1 < 2/\mu_2$, then the expected waiting time k/μ_k is strictly increasing in the initial queue position k, by Lemma 1. If the information rule I admits two distinct beliefs at $t = 0$ ordered in FOSD, then the associated expected waiting times are distinct, so the result again follows from Proposition 3. ∎

6.2 SIRO

In SIRO, agents are served with equal probability up to the service capacity. Specifically, under SIRO, each agent in a queue is served at rate $q_{k,\ell} = \mu_k/k$, regardless of his queue

[29] Two beliefs, γ and γ' are ordered in FOSD if $\sum_{i=1}^{k} \gamma'_j \leq \sum_{i=1}^{k} \gamma_j$ for all $k = 1, ..., K^*$, or vice versa.

position $\ell \leq k$. The expected waiting time can then be characterized as a function of the current length of the queue.

Lemma 4. (Waiting time under SIRO) Let τ_k be the expected waiting time under SIRO for an agent who enters a queue with $k-1$ agents. The tuple $\tau = (\tau_1, ..., \tau_{K^*})$ is well defined and τ_k is nondecreasing in k. Further, $\tau_1 \geq 1/\mu_1$ and $\tau_{K^*} \leq K^*/\mu_{K^*}$. If $1/\mu_1 < K^*/\mu_{K^*}$, then $\tau_1 < \tau_{K^*}$. If $1/\mu_1 < 2/\mu_2$, then τ_k is strictly increasing in k.

Proof. See Appendix D.3. ∎

☐ **SIRO with information.** Lemma 4 enables us to produce the same corollary for SIRO as for FCFS, suggesting that providing agents with nontrivial information about the queue length is generally suboptimal.

Corollary 2 (SIRO with information)**.** If $1/\mu_1 < K^*/\mu_{K^*}$ then SIRO with full information cannot implement the cutoff policy. If $1/\mu_1 < 2/\mu_2$ and information rule I admits $\gamma, \gamma' \in \text{supp}(I_0)$, where $\gamma' \neq \gamma$ are ordered in FOSD, then SIRO with I cannot implement the optimal policy.

Proof. The arguments are analogous to the proof of Corollary 1, upon using Lemma 4. ∎

☐ **SIRO with no information.** We next consider SIRO when no information is provided to agents beyond the recommendations to join the queue or to stay in it. To investigate the implementability of the optimal outcome with this queue/information design, we must examine the dynamic evolution of agents' beliefs conditional on staying in the queue, again characterized via a system of ODEs (see Appendix D.4).

Proposition 4 (SIRO with no information)**.** Assume the service process is regular, $1/\mu_1 < K^*/\mu_{K^*}$, and $\lambda_k \left(1 - \frac{\mu_{k+1}}{\mu_k} \frac{k}{k+1}\right)$ is strictly decreasing in k for $k = 2, ..., K^* - 1$. Then, SIRO cannot implement the optimal cutoff policy under *no information*—hence under *any* information policy.

Proof. See Appendix D.4. ∎

One can see that the conditions of Proposition 4 hold in the most canonical setting:

Corollary 3. SIRO with no information cannot implement the optimal cutoff policy in the $M/M/1$ queue.

Proof. In the $M/M/1$ queue, $\mu_k = \mu$ and $\lambda_k = \lambda$ for all k, for some $\mu, \lambda > 0$. Hence, $1/\mu < K^*/\mu$ for $K^* > 1$. Further, $\lambda_k \left(1 - \frac{\mu_{k+1}}{\mu_k} \frac{k}{k+1}\right) = \lambda \frac{1}{k+1}$, which is strictly decreasing in k. Hence, the result follows from Proposition 4. ∎

7 Necessity of FCFS for Optimality in a Rich Domain

In principle, there can be other queue designs. Although unobserved in practice, queueing disciplines such as LCFS and LIEW (Load Independent Expected Wait) have been studied in the literature. Instead of studying these queueing disciplines separately, we consider *all* feasible queueing disciplines and all (regular) environments and show that FCFS is the only queueing rule that is optimal for all (regular) queueing environments. In other words, for any queueing discipline differing from FCFS, there is a (regular) environment in which it cannot implement the optimal cutoff policy solving $[P']$ under *any* information policy. This will prove the *necessity* of FCFS to implement the solution of $[P']$ in a rich domain.

Since we have already established that FCFS with no information can implement the solution of $[P']$ in any regular environment, it suffices to show that for each discipline differing from FCFS, there exists *some* regular environment under which that mechanism cannot implement the optimal cutoff policy. To this end, we focus on a $M/M/1$ queue environment in which (i) the optimal cutoff policy involves the maximal length $K^* = 2$, (ii) no rationing at $k = 1$ (i.e., $x_1^* = 1$), and (iii) (IR) is binding at the optimal cutoff policy that solves $[P']$.[30]

Given $K^* = 2$, a queue discipline matters only when there are two agents in the queue. In that case, let π denote the probability that the first agent receives first priority, or equivalently, that newly arrived agents receive second priority. Then, a newly arrived agent receives first priority with probability $1 - \pi$. Clearly, FCFS corresponds to $\pi = 1$, SIRO corresponds to $\pi = 1/2$, and LCFS corresponds to $\pi = 0$. LIEW is designed to ensure that an agent faces the same expected waiting time when he joins the queue regardless of the current queue length. It will be seen that $\pi = \frac{\lambda}{2\lambda+p} \in (0, 1/2)$ implements this rule. It is not difficult to see that any feasible Markovian queueing discipline is characterized fully by parameter π.

Fix any $\pi \in [0, 1]$. There are only three relevant "states," $(k, \ell) = (1, 1), (2, 1), (2, 2)$, based on the length k of the queue and one's service priority ℓ. Let $\tau_{1,1}, \tau_{2,1}$ and $\tau_{2,2}$ denote the expected waiting times for an agent, when he is the only one in the queue, when he has first priority among two agents, and when he has second priority among two agents, respectively. The expected waiting time $\tau_{1,1}$ must satisfy:

$$\tau_{1,1} = (\mu dt)dt + \lambda dt(dt + \pi\tau_{2,1} + (1 - \pi)\tau_{2,2}) + (1 - \lambda dt - \mu dt)(dt + \tau_{1,1}) + o(dt),$$

since, for a small time increment dt, the sole agent in the queue waits for time dt if he is served during that period (which occurs with probability μdt), for $dt + \pi\tau_{2,1} + (1 - \pi)\tau_{2,2}$ if another agent arrives (which occurs with probability λdt), and for $dt + \tau_{1,1}$ if neither event arises (which occurs with probability $1 - \lambda dt - \mu dt$). Similar reasoning yields:

$$\tau_{2,1} = (\mu dt)dt + (1 - \mu dt)(dt + \tau_{1,1}) + o(dt)$$
$$\tau_{2,2} = (\mu dt)(dt + \tau_{1,1}) + (1 - \mu dt)(dt + \tau_{2,2}) + o(dt).$$

[30]This is the uniquely optimal policy if $V/C = \frac{2\lambda+\mu}{(\lambda+\mu)\mu}$ and $\alpha = 0$.

Letting $dt \to 0$ and simplifying, we obtain:

$$\tau_{1,1} = \frac{\mu + \lambda}{\mu(\mu + \lambda\pi)}, \tau_{2,1} = \frac{1}{\mu}, \tau_{2,2} = \frac{\mu + \lambda}{\mu(\mu + \lambda\pi)} + \frac{1}{\mu}. \tag{3}$$

One can see that

$$\tau_{2,1} \leq \tau_{1,1} < \tau_{2,2}.$$

The intuition for the first inequality, which holds with equality only when $\pi = 1$ (i.e., FCFS), is that an agent who has service priority in a $k = 2$ queue never loses that priority until he is served (since there is no more entry), whereas an agent in a $k = 1$ queue may lose his priority if somebody else enters the queue before he is served.

When an agent enters a $k = 1$ queue, he expects to wait for $\pi\tau_{2,2} + (1 - \pi)\tau_{2,1}$. When he enters an empty queue, then he expects to wait for $\tau_{1,1}$. It is intuitive, and one can easily confirm, that $\pi\tau_{2,2} + (1 - \pi)\tau_{2,1} > \tau_{1,1}$ under FCFS (i.e., when $\pi = 1$) and under SIRO (i.e., when $\pi = 1/2$). It is also intuitive that the opposite is true, namely $\pi\tau_{2,2} + (1-\pi)\tau_{2,1} < \tau_{1,1}$, under LCFS (i.e., when $\pi = 0$). Equalizing the two waiting times gives rise to LIEW, or $\pi = \frac{\lambda}{2\lambda + \mu} \in (0, 1/2)$.

We now consider the information policy that provides no information (beyond the recommendations) for all $t \geq 0$. This is without loss since, if a queueing rule π fails (IC_t), for some $t \geq 0$, under no information, it would fails (IC_t) under *any* information policy. For our purposes, we study the evolution of an agent's beliefs $\gamma^t = (\gamma^t_{1,1}, \gamma^t_{2,1}, \gamma^t_{2,2})$ on the relevant state (k, ℓ) over time, where $k = 1, 2, \ell \leq k$. Under no information, there is a unique belief γ^t in $\text{supp}(I_t)$ at each $t \geq 0$. Of particular interest is the expected waiting time conditional on having waited for $t \geq 0$ in the queue:

$$\omega(t) = \gamma^t_{1,1}\tau_{1,1} + \gamma^t_{2,1}\tau_{2,1} + \gamma^t_{2,2}\tau_{2,2}.$$

Although suppressed for notational ease, both $\gamma^t = (\gamma^t_{1,1}, \gamma^t_{2,1}, \gamma^t_{2,2})$ and $\tau = (\tau_{1,1}, \tau_{2,1}, \tau_{2,2})$ depend on π.

The next result provides a crucial characterization:

Proposition 5. If $\pi < \frac{\mu+\lambda}{\mu+2\lambda}$, then there exists $t > 0$ such that $\omega(t) > \omega(0)$, and π fails (IC_t) under no information (again under the assmption that (IR) binds at the optimal cutoff policy).

Proof. See Appendix E. ∎

Figure 2 illustrates expected waiting times under different values of π corresponding to LCFS, LIEW, SIRO, and FCFS. As is clearly seen, as time passes an agent in the queue expects to wait increasingly more under the first three rules and increasingly less under FCFS.

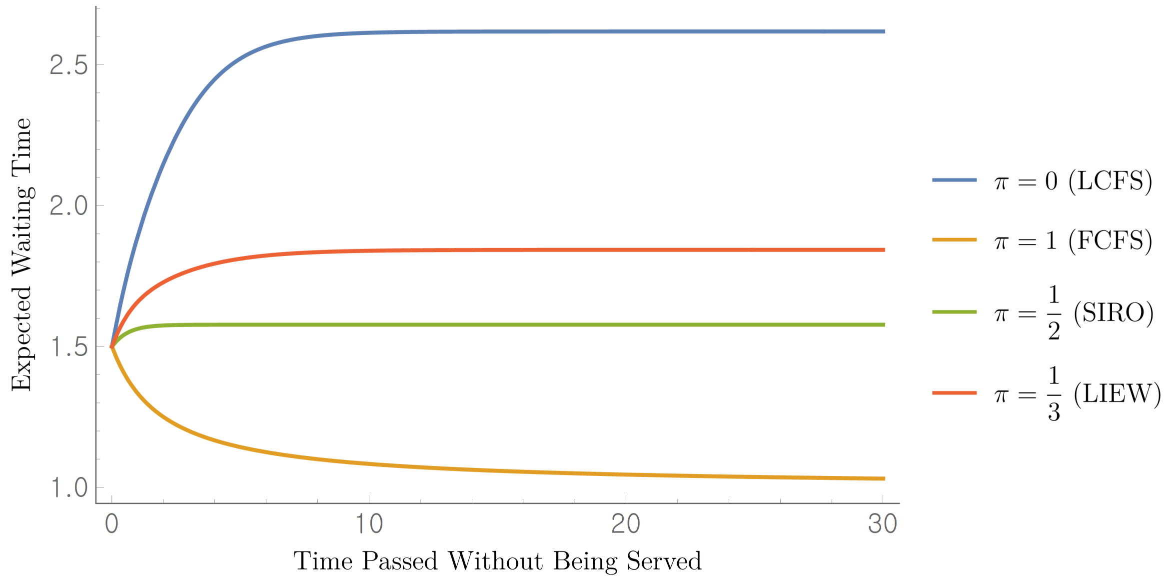

Figure 2: Expected waiting times under alternative values of π.
Note: $M/M/1$ with $K^* = 2$; $\lambda = \mu = 1$.

To see why a low π leads to less favorable belief as time passes, it is instructive to consider LCFS with $\pi = 0$. Panel (a) of Figure 3 depicts how an agent's beliefs over states $(k, \ell) = (2, 1), (1, 1), (2, 2)$ evolve over time. Under LCFS, when an agent joins a queue (i.e., at $t = 0$), he can be sure that he is not at state $(k, \ell) = (2, 2)$; either the queue was empty, in which case he is at $(1, 1)$ or else, he receives priority, so $(k, \ell) = (2, 1)$. In the latter case, he is likely to be served soon, whereas the former means that he will not be served as quickly. As time passes, the fact that he is still in the queue without having been served gives more credence to the former possibility $(k, \ell) = (1, 1)$. This is why as time passes, the updating becomes more unfavorable, with $\gamma^t_{1,1}$ rising faster than $\gamma^t_{2,1}$.

Like LCFS, LIEW has $\pi < 1/2$, and a similar intuition applies; as time passes without getting served, an agent's belief that $(k, \ell) = (2, 1)$ updates downward; see Panel (b) of Figure 3.[31]

We are now ready to state our main result:

Theorem 3. For any $\pi < 1$, there exists (V, C, λ, μ) such that the queueing rule π fails (IC_t) for some $t > 0$ under *any* information. Hence, π cannot implement the optimal cutoff policy under any information policy.

Proof. For any π, there exist λ and μ such that $\pi < \frac{\mu+\lambda}{\mu+2\lambda}$. Next, for V and C such that $V = \frac{2\lambda+\mu}{(\lambda+\mu)\mu} C$, under the optimal policy, (IR) binds with $K^* = 2$ and there is no rationing

[31] As pointed out earlier, similar unfavorable updating on the initial state occurs with FCFS, but the belief on $(k, \ell) = (2, 2)$ falls over time and this favorable belief evolution dominates the unfavorable updating on the initial queue length. See panel (c) of Figure 3.

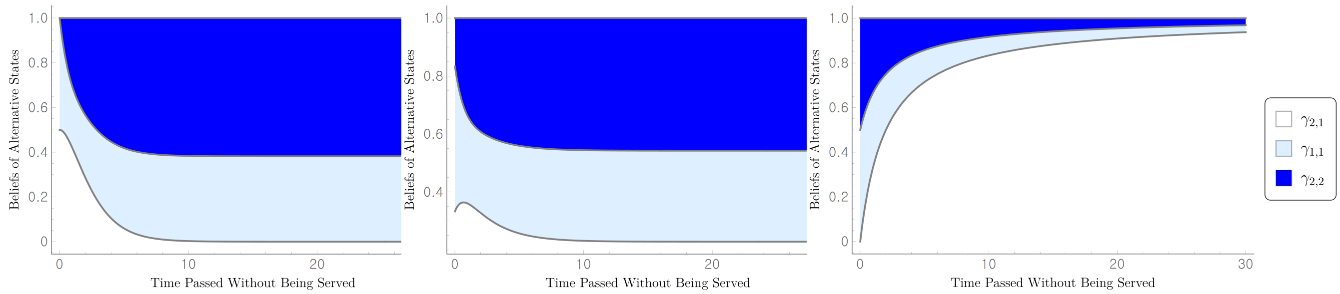

(a) $\pi = 0$ (LCFS) (b) $\pi = 1/3$ (LIEW) (c) $\pi = 1$ (FCFS)

Figure 3: Evolution of beliefs under alternative queueing disciplines
Note: $M/M/1$ with $K^* = 2$; $\lambda = \mu = 1$.

at $k = 1$. By Proposition 5, the rule π fails (IC_t) for some $t > 0$ under a "no" information rule. This in turn implies that for the same t, (IC_t) must fail under *any* information. This is due to the martingale property, if γ^t fails (IC_t) under no information, for any other I, we must have some $\tilde{\gamma}^t \in \text{supp}(I_t)$ that must fail (IC_t). Hence, π fails to implement the optimal solution to $[P']$ under any information policy. ∎

8 Concluding Remarks

We have focused on a canonical queueing model involving a single queue. But the insights we obtain appear general and apply beyond our model. Here we discuss how one may extend our analysis to other settings of potential interest.

Dynamic two-sided matching. As noted earlier, there is a growing interest in extending two-sided matching analysis to a dynamic framework. Leshno (2019) and Baccara, Lee, and Yariv (2020) consider models in which two different types of agents match with either two different types of objects (e.g., housing) or two different types of agents. In such a model, efficiency demands that the designer accumulates agents in a queue until a right type of object or agent arrives, to avoid mismatching. Leshno (2019) assumes overloaded demand so that the planner wishes to incentivize the agents to queue as much as possible, and demonstrates in the full information model that SIRO outperforms FCFS in this regard, and LIEW, by equalizing the waiting time regardless of the queue length, outperforms all other mechanisms. Despite the ostensible difference relative to our model, Section S.4 in the online appendix shows that our analysis applies without much modification to this model, and points out that the main results from Leshno (2019) rest crucially on his full information assumption. With optimal information design, the FCFS could do just as well as any other mechanism, including LIEW, in incentivizing agents to enter a queue. If one includes the dynamic incentive problem, which Leshno (2019) does not consider,[32] then FCFS does strictly better than other queueing disciplines. Baccara, Lee, and Yariv (2020)'s model is similar to that of Leshno (2019), except that there are agents on both sides. Hence, our main insight of Theorem 2 applies, except for one difference. Unlike Leshno (2019), agents' incentives to enter a queue may be excessive under FCFS with full information. While this is an issue in their decentralized matching, in our setting the designer can easily solve the problem by preventing an agent from entering a queue, as is often done in practice. Meanwhile, they also show that a queueing discipline admits insufficient entry under LCFS. In that case, information design can be useful even with their assumption. In addition, they too do not consider dynamic incentives, for which our analysis will prove useful.

Transfers. We have assumed that the designer is unable to use transfers, say to incentivize agents. This assumption is based on realism; monetary compensation is rarely used in practice to incentivize queueing. One difficulty with transfers may be to discern genuine customers from "pretenders" who may not have real demand for service; if a monetary reward is given out, a customer may call a customer service even without having a genuine issue with a product. Setting aside such difficulties, it is rather straightforward to extend our analysis to allow for transfers. With transfers, $[P']$ reduces to the maximization of a sum of W and agents' utilities without (IR),[33] and the ensuring problem is qualitatively similar to that of $[P']$. In particular, both Theorem 1 and Theorem 2 remain valid. Intuitively, the fact that FCFS with no information provides better (dynamic) incentives than other policies means that even when (IR) binds so that transfers are required to subsidize agents' entry, the policy will minimize the amount of entry subsidy and therefore will continue to be desirable. In short, the main thrust of our analysis and results will generalize to this environment.

Avenues of Future Research. There are several interesting avenues along which one could extend the current model. First, it would be interesting to relax the homogeneity of agents' preferences. In practice, it is reasonable that agents differ in their waiting costs

[32] The dynamic incentive issue does not arise in SIRO or FCFS under complete information: any agent who joins the queue will have the incentive to stay in the queue. Recall, however, that neither discipline would implement the optimum under complete information. Under *no information* (which is optimal), dynamic incentives will be an issue. Although an agent may not leave the queue and unilaterally "claim" a mismatched object, which is presumably under the designer's control, he/she may leave the queue without claiming any object. If the value of outright exit is not very low (e.g., in comparison with the value of a mismatched object), then the dynamic incentives will matter just as they do in our model. Specifically, both SIRO (under no information) and LIEW (under *any* information) would be vulnerable to reneging, and cannot implement the optimal outcome, as stated in our Theorem 3.

[33] This can be seen from substituting the constraint (now with transfer) into the objective function when (IR) is binding. The logic is analogous to the well-known fact that Pareto optimality reduces to utilitarian welfare maximization when there are perfect transfers.

and/or in their value of service. The optimal policy will then call for treating agents differently based on their heterogeneous preferences—for instance, encouraging agents with low waiting costs to join the queue and discouraging those with high waiting costs from doing so. To the extent that preference heterogeneity is unlikely to be observed, eliciting preferences will also become important. How the tradeoff between these two considerations must be balanced is an interesting problem that warrants investigation. Another interesting avenue of research is to consider uncertainty in the agents' arrival or server's capacity. Such an uncertainty about the primitive process will create incentives for agents to learn about the current "state," and will engender a nontrivial dynamic consideration in their decision making.[34] These and other worthy extensions of the current model await future research.

References

Akbarpour, M., J. Combe, V. Hiller, R. Shimer, and O. Tercieux (2020): "Unpaired Kidney Exchange: Overcoming Double Coincidence of Wants without Money," *NBER Working paper number 27765.* 7

Akbarpour, M., S. Li, and S. O. Gharan (2020): "Thickness and information in dynamic matching markets," *Journal of Political Economy*, 128(3), 783–815. 6, 13

Anunrojwong, J., K. Iyer, and V. Manshadi (2020): "Information Design for Congested Social Services: Optimal Need-Based Persuasion," *EC '20: Proceedings of the 21st ACM Conference on Economics and Computation*, pp. 349–350. 7

Ashlagi, I., M. Burq, P. Jaillet, and V. Manshadi (2019): "On Matching and Thickness in Heterogeneous Dynamic Market," *Operations Research*, 67(4), 927–949. 7

Ashlagi, I., M. Faidra, and A. Nikzad (2020): "Optimal Dynamic Allocation: Simplicity through Information Design," Discussion paper, Stanford. 7

Baccara, M., S. Lee, and L. Yariv (2020): "Optimal dynamic matching," *Theoretical Economics*, 15, 1221–1278. 7, 13, 29, 30

Bloch, F., and D. Cantala (2017): "Dynamic assignment of objects to queuing agents," *American Economic Journal: Microeconomics*, 9, 88–122. 6

Cripps, M. W., and C. D. Thomas (2019): "Strategic experimentation in queues," *Theoretical Economics*, 14, 647–708. 5, 31

Doval, L., and B. Szentes (2018): "On the efficiency of queuing in dynamic matching," Discussion paper, Caltech and London School of Economics. 7, 13

[34]See Cripps and Thomas (2019) for such a model.

Gnedenko, B., and I. Kovalenko (1989): *Introduction to Queueing Theory*. Birkhauser. 12

Hassin, R. (1985): "On the optimality of first come last served queues," *Econometrica*, 53, 201–202. 2, 6, 12

Hassin, R. (2012): *Rational Queueing*. CRC Press. 5

Hassin, R., and M. Haviv (1995): "Equilibrium strategies for queues with impatient customers," *Operations Research Letters*, 17, 41–45. 5, 12, 21

——— (2012): *To Queue or Not to Queue: Equilibrium Behavior in Queueing Systems*. Kluwer Academic Publishers. 5, 6, 11

Haviv, M., and B. Oz (2016): "Regulating an observable M/M/1 queue," *Operations Research Letters*, 44, 196–198. 6

——— (2018): "Self-Regulation of an Unobservable Queue," *Management Science*, 64, 1–10. 6

Haviv, M., and Y. Ritov (2001): "Homogeneous customers renege from invisible queues at random times under deteriorating waiting conditions," *Queueing Systems*, 38, 495–508. 5, 12, 21

Kamenica, E., and M. Gentzkow (2011): "Bayesian persuasion," *American Economic Review*, 101, 2590–2615. 9

Kipp, M., C. Ryan, and S. Matt (2016): "The Slater Conundrum: Duality and Pricing in Infinite-Dimensional Optimization," *SIAM Journal on Optimization*, 1(26), 111–138. 34

Kittsteiner, T., and B. Moldovanu (2005): "Priority auctions and queue disciplines that depend on processing time," *Management Science*, 51, 236–248. 6

Larson, R. C. (1987): "Perspectives on queues: social justice and the psychology of queueing," *Operations Research*, 35, 895–905. 7

Leshno, J. (2019): "Dynamic matching in overloaded waiting lists," Discussion paper, SSRN Working Paper 2967011. 6, 13, 29, 30

Lingenbrink, D., and K. Iyer (2019): "Optimal signaling mechanisms in unobservable queues," *Operations Research*, 67, 1397–1416. 7, 12, 16

Mandelbaum, A., and N. Shimkin (2000): "A model for rational abandonments from invisible queues," *Queueing Systems*, 36, 141–173. 5, 9

Margaria, C. (2020): "Queueing to learn," Discussion paper, Boston University. 6

Myerson, R. B. (1981): "Optimal auction design," *Mathematics of Operations Research*, 6(1), 58–73. 7

Naor, P. (1969): "The regulation of queue size by levying tolls," *Econometrica*, 37, 15–24. 2, 5, 6, 12, 23

Platz, T. T., and L. P. Østerdal (2017): "The curse of the first-in–first-out queue discipline," *Games and Economic Behavior*, 104, 165–176. 2, 6

Shaler Stidham, J. (2009): *Optimal Design of Queueing Systems*. CRC Press. 7

Sherzer, E., and Y. Kerner (2018): "Customers' abandonment strategy in an M/G/1 queue," *Queueing Systems*, 90, 65–87. 5

Simhon, E., Y. Hayel, D. Starobinski, and Q. Zhu (2016): "Optimal information disclosure policies in strategic queueing games," *Operations Research Letters*, 44, 1109–113. 7, 12

Su, X., and S. Zenios (2004): "Patient choice in kidney allocation: The role of the queueing discipline," *Manufacturing and Services Operations Management*, 6, 280–301. 2, 6

A Proof of Theorem 1

We first denote by M the set of distributions in $\Delta(\mathbb{Z}_+) : p = \{p_k\}_{k=0}^{\infty}$ satisfying the following constraint

$$\lambda_k p_k - \mu_{k+1} p_{k+1} \geq 0, \forall k = 0, ...$$

We rewrite the LP problem $[P']$ as:

$$[P'] \qquad \max_{p \in M} \sum_{k=0}^{\infty} p_k \left[\mu_k((1-\alpha)R + \alpha V) - \alpha C k\right]$$

subject to

$$\sum_{k=0}^{\infty} p_k \left[\mu_k V - Ck\right] \geq 0.$$

(recall our convention that $\mu_0 = 0$). While this is an LP problem, it is infinite dimensional. In particular, a saddle point characterization of the optimal solution, which we use, may not be assured.[A.35] Since this is a critical part of our argument below, we first study a finite dimensional truncation of $[P']$ where the state space contains finitely many states, say K, where K can potentially be "large". In this environment, we will show that an optimal solution p^K exhibits a cutoff policy (Appendix A.1). In a second step, we will show that as K gets large, a limit point of $\{p^K\}$ is an optimal solution of $[P']$ and exhibits a cutoff policy (Appendix A.2).

A.1 Finite dimensional analysis

We denote by M_K the set of distributions $p = \{p_k\}_{k=0}^{K}$ in $\Delta(\{0, 1, ..., K\})$ satisfying the following constraints

$$\lambda_k p_k - \mu_{k+1} p_{k+1} \geq 0, \forall k = 0, ..., K-1.$$

We consider the following "truncated" version of $[P']$, say $[P'_K]$

$$[P'_K] \qquad \max_{p \in M_K} \sum_{k=0}^{K} p_k \left[\mu_k((1-\alpha)R + \alpha V) - \alpha C k\right]$$

subject to

$$\sum_{k=0}^{K} p_k \left[\mu_k V - Ck\right] \geq 0.$$

[A.35] Countably infinite linear programs (CILPs) are linear optimization problems with a countably infinite number of variables and a countably infinite number of constraints. It is well-known that many of the nice properties of finite dimensional linear programming may fail to hold in these problems. Indeed, while in finite dimensional LP problems, Slater's condition—ensuring zero duality gap— is very weak, infinite dimensional versions of Slater's condition are much more demanding and zero duality gap may often fail. See Kipp, Ryan, and Matt (2016) and references therein.

Let us fix $\xi \geq 0$ and consider the problem $[\mathcal{L}_\xi]$

$$[\mathcal{L}_\xi] \qquad \max_{p \in M_K} \mathcal{L}(p, \xi)$$

where

$$\mathcal{L}(p, \xi) \triangleq \sum_{k=0}^{K} p_k \left[\mu_k((1-\alpha) R + \alpha V) - \alpha C k\right] + \xi \sum_{k=0}^{K} p_k \left[\mu_k V - C k\right].$$

We can write $\mathcal{L}(p, \xi)$ as

$$\sum_{k=0}^{K} p_k f(k; \xi)$$

where $f(k; \xi) \triangleq \mu_k((1-\alpha)R + (\alpha+\xi) V) - (\alpha+\xi) C k$.

The Lagrangian dual of problem $[P'_K]$ is taking the inf over $\xi \geq 0$ of the value of $[\mathcal{L}_\xi]$. Since M_K is a convex set, the problem constitutes a finite dimensional linear program, so strong duality applies. In such environments, p^* is an optimal solution if and only if there is (a Lagrange multiplier) $\xi^* \geq 0$ such that (p^*, ξ^*) is a saddle point of the function $\mathcal{L}(\cdot, \cdot)$, i.e.,

$$\mathcal{L}(p, \xi^*) \leq \mathcal{L}(p^*, \xi^*) \leq \mathcal{L}(p^*, \xi)$$

for any $\xi \geq 0$ and $p \in M_K$. We use this saddle point characterization in the remaining parts of our proof. In the sequel, we fix (p^*, ξ^*) a saddle point of the function $\mathcal{L}(\cdot, \cdot)$.

In this section, we will show a finite-dimensional version of Theorem 1 stated below.

Proposition A.6. If μ is regular, then there is an optimal solution for $[P'_K]$ which exhibits a cutoff policy. In addition, $p^*_k > 0$ for each $k \leq \min\{k^*, K\}$ where $k^* \triangleq \min \arg\max f(k; \xi^*)$.

In the sequel, we say that a function $f : \mathbb{Z}_+ \to \mathbb{R}$ is *single-peaked* if $f(k-1) < f(k)$ for all $k \leq \min \arg\max_{k \in \mathbb{Z}_+} f(k)$ while $f(k) > f(k+1)$ for all $k \geq \max \arg\max_{k \in \mathbb{Z}_+} f(k)$. Our convention is that if $\arg\max_{k \in \mathbb{Z}_+} f(k)$ is empty, then $\min \arg\max_{k \in \mathbb{Z}_+} f(k)$ is set to $+\infty$. We now show that regularity implies single-peakedness of $f(\cdot; \xi)$.

Lemma A.5. If μ is regular, then for any $\xi \geq 0$, function $f(\cdot; \xi)$ is single-peaked.

Proof. Fix any $\xi \geq 0$. We start by showing that if $f(k; \xi) \geq f(k+1; \xi)$ then $f(k'; \xi) \geq f(k'+1; \xi)$ for any $k' \geq k$, i.e., if the function decreases at some state k, then it decreases at larger states. Indeed, assume that $f(k; \xi) \geq f(k+1; \xi)$, i.e.,

$$\mu_k((1-\alpha) R + (\alpha+\xi) V) - (\alpha+\xi) C k \geq \mu_{k+1}((1-\alpha) R + (\alpha+\xi) V) - (\alpha+\xi) C (k+1).$$

Simple algebra shows that this is equivalent to

$$\mu_{k+1} - \mu_k \leq \frac{(\alpha+\xi) C}{(1-\alpha) R + (\alpha+\xi) V}.$$

Since μ is regular, $\mu_{k+1} - \mu_k$ is nonincreasing and so, for $k' \geq k$, we must have

$$\mu_{k'+1} - \mu_{k'} \leq \mu_{k+1} - \mu_k \leq \frac{(\alpha+\xi)C}{(1-\alpha)\,R + (\alpha+\xi)V}.$$

Hence, $f(k';\xi) \geq f(k'+1;\xi)$, as claimed.

Now, let $k^* \triangleq \min \arg\max_{k\in\mathbb{Z}_+} f(k;\xi)$; recall that our convention allows k^* to be $+\infty$. Now, there are two candidates. First, $k^* = +\infty$. In that case, since we showed that whenever the function $f(\cdot;\xi)$ decreases at some state k, it must decrease at larger states, it must be that $f(\cdot;\xi)$ is a strictly increasing function and we are done. Second, $k^* < +\infty$. Fix $k \leq k^*$. We need to show that

$$\mu_{k-1}((1-\alpha)\,R + (\alpha+\xi)V) - (\alpha+\xi)C(k-1) < \mu_k((1-\alpha)\,R + (\alpha+\xi)V) - (\alpha+\xi)Ck.$$

Simple algebra shows that this is equivalent to

$$\mu_k - \mu_{k-1} > \frac{(\alpha+\xi)C}{(1-\alpha)\,R + (\alpha+\xi)V}.$$

By definition of k^*, function $f(\cdot;\xi)$ achieves a maximum at k^* and k^*-1 is not a maximum, hence,

$$\mu_{k^*} - \mu_{k^*-1} > \frac{(\alpha+\xi)C}{(1-\alpha)\,R + (\alpha+\xi)V}.$$

Since μ is regular, $\mu_k - \mu_{k-1}$ is nonincreasing and so, for $k \leq k^*$, we must have

$$\mu_k - \mu_{k-1} \geq \mu_{k^*} - \mu_{k^*-1} > \frac{(\alpha+\xi)C}{(1-\alpha)\,R + (\alpha+\xi)V}.$$

It thus follows that for any $k \leq k^*$,

$$\mu_{k-1}((1-\alpha)\,R + (\alpha+\xi)V) - (\alpha+\xi)C(k-1) < \mu_k((1-\alpha)\,R + (\alpha+\xi)V) - (\alpha+\xi)Ck,$$

as claimed. The argument for $k \geq \max \arg\max_{k\in\mathbb{Z}_+} f(k;\xi)$ is similar. ∎

Before proceeding, we make the following straightforward observations (1) $p_0^* > 0$ (or else $p_k^* = 0$ for all k); (2) for all ξ, $f(0;\xi) = 0$. Using these two facts, we claim that Proposition A.6 holds whenever $f(k;\xi^*) = f(k';\xi^*)$ for all k, k' in the support of p^*. Indeed, since $p_0^* > 0$, $f(k;\xi^*) = 0$ for all states k in the support of p^*. In that case, $\sup_p \mathcal{L}(p,\xi^*) = 0$. Thus, the value of the problem $[P_K']$ is 0. Clearly, the distribution p corresponding to the Dirac measure on state 0 yields the same value and is a cutoff policy. Hence, in this very special case, Theorem 1 holds true. Thus, in the sequel, we assume that there is a pair of states k and k' in the support of p^* satisfying $f(k;\xi^*) \neq f(k';\xi^*)$. We will use the following claim.

Lemma A.6. Fix any $\ell \leq K-1$ such that

$$f(\ell;\xi^*) < f(\ell+1;\xi^*).$$

We must have $\mu_{\ell+1}p^*_{\ell+1} = \lambda_\ell p^*_\ell$.

Proof. Fix ℓ satisfying the properties of the claim. Since p^* is an optimal solution of $[P'_K]$, we know that $\mu_{\ell+1}p^*_{\ell+1} \leq \lambda_\ell p^*_\ell$. Assume that $\mu_{\ell+1}p^*_{\ell+1} < \lambda_\ell p^*_\ell$ (otherwise, we are done). Now, simply consider $\hat{p}$ which is defined by $\hat{p}_{\ell+1} = p^*_{\ell+1}+\varepsilon$ and $\hat{p}_\ell = p^*_\ell-\varepsilon$ ($\hat{p}$ is equal to p^* otherwise) and note that we can choose $\varepsilon > 0$ so that $\mu_{\ell+1}\hat{p}_{\ell+1} = \lambda_\ell\hat{p}_\ell$ while ensuring $\hat{p}_\ell, \hat{p}_{\ell+1} \in [0,1]$.[A.36] Clearly, $\sum_{k=0}^K \hat{p}_k = 1$. Now, let us show that $\mu_{k+1}\hat{p}_{k+1} \leq \lambda_k\hat{p}_k, \forall k = 0, ...K-1$. Since this constraint holds when replacing $\hat{p}$ by p^* (because p^* is an optimal solution of $[P'_K]$), by construction of $\hat{p}$, we only need to check this constraint for $k = \ell+1$ and $k = \ell-1$. For the constraint at $k = \ell+1$, note that since the constraints of the LP problem hold for p^*, we have

$$\mu_{\ell+2}\hat{p}_{\ell+2} = \mu_{\ell+2}p^*_{\ell+2} \leq \lambda_{\ell+1}p^*_{\ell+1} \leq \lambda_{\ell+1}\hat{p}_{\ell+1}.$$

Similarly, for the constraint at $k = \ell-1$,

$$\mu_\ell\hat{p}_\ell \leq \mu_\ell p^*_\ell \leq \lambda_{\ell-1}p^*_{\ell-1} = \lambda_{\ell-1}\hat{p}_{\ell-1}.$$

Now, we show that the value of the objective of $[\mathcal{L}_{\xi^*}]$ strictly increases when we replace solution p^* by $\hat{p}$. We have

$$\begin{aligned}\sum_{k=0}^K \hat{p}_k f(k;\xi^*) - \sum_{k=0}^K p^*_k f(k;\xi^*) &= \hat{p}_\ell f(\ell;\xi^*) - p^*_\ell f(\ell;\xi^*) + \hat{p}_{\ell+1} f(\ell+1;\xi^*) - p^*_{\ell+1} f(\ell+1;\xi^*) \\ &= -\varepsilon f(\ell;\xi^*) + \varepsilon f(\ell+1;\xi^*) = \varepsilon\left(f(\ell+1;\xi^*) - f(\ell;\xi^*)\right) > 0\end{aligned}$$

where the inequality comes from the assumption in the claim. To conclude, we must have that $\mathcal{L}(\hat{p},\xi^*) > \mathcal{L}(p^*,\xi^*)$ which contradicts the fact that (p^*,ξ^*) is a saddle point of the function $\mathcal{L}(\cdot,\cdot)$. ∎

In the proof of Proposition A.6, we will need the following simple lemma which proof is relegated to Section S.1 of the online appendix.

Lemma A.7. Assume that p' stochastically dominates p. Let φ be a nondecreasing function. If there is κ such that

$$\sum_{k=\kappa}^K p'_k > \sum_{k=\kappa}^K p_k$$

[A.36] Indeed, at $\varepsilon = 0$, we have $\mu_{\ell+1}\hat{p}_{\ell+1} < \lambda_\ell\hat{p}_\ell$. In addition, for $\varepsilon = p_\ell > 0$ we have $\hat{p}_{\ell+1} = p_{\ell+1} + \varepsilon = p_{\ell+1} + p_\ell \leq 1$ and $\mu_{\ell+1}\hat{p}_{\ell+1} > \lambda_\ell\hat{p}_\ell = 0$. Hence, by the Intermediate Value Theorem, there must exist $\varepsilon \in (0, p_\ell)$ so that $\mu_{\ell+1}\hat{p}_{\ell+1} = \lambda_\ell\hat{p}_\ell$.

and $\varphi(\kappa) > \varphi(\kappa - 1)$ then

$$\sum_{k=0}^{K} p'_k \varphi(k) > \sum_{k=0}^{K} p_k \varphi(k).$$

Proof of Proposition A.6. Let k^* be $\min \arg\max_k f(k;\xi^*)$ and k^{**} be $\max \arg\max_k f(k;\xi^*)$. Recall that k^* can be equal to $+\infty$. By Lemma A.6, we know that $f(k;\xi^*)$ is strictly increasing up to k^*. Hence, the above claim implies that $\mu_k p^*_k = \lambda_{k-1} p^*_{k-1}$ for each $k \leq \min\{k^*, K\}$. Note that this also implies that $p^*_k > 0$ for each $k \leq \min\{k^*, K\}$, as stated in Proposition A.6. If $K \leq k^*$, we are done. So assume from now on that $K > k^*$; note that this implies that $k^* < +\infty$. By means of contradiction, let us assume that p^* does not exhibit a cutoff policy. This means that there is $k_0 > k^*$ at which $\mu_{k_0} p^*_{k_0} < \lambda_{k_0-1} p^*_{k_0-1}$ and $p^*_{k_0+1} > 0$ (hence, $p^*_{k_0} > 0$). Wlog, assume that for any $k < k_0$, we have $\mu_k p^*_k = \lambda_{k-1} p^*_{k-1}$. We consider two cases.

Case 1 : $p^*_k > 0$ for some $k > k^{**}$. We build $\hat{p}$ which strictly increases the value of the objective of $[\mathcal{L}_{\xi^*}]$ over p^*. Define $\hat{p}$ to be equal to p^* for $k \leq k_0 - 1$. Then, for each $k \geq k_0$, build $\hat{p}$ inductively so that $\mu_k \hat{p}_k = \lambda_{k-1} \hat{p}_{k-1}$. Since the total mass of $\hat{p}$ must be equal to 1, this may be possible only up to a point. Hence, there is $\hat{K} \geq k_0$ (potentially equal to K) such that $\mu_k \hat{p}_k = \lambda_{k-1} \hat{p}_{k-1}$ for $k = 0, ..., \hat{K} - 1$, and $\hat{p}_k = 0$ for $k > \hat{K}$. One can show inductively that $\hat{p}_k > p^*_k$ for all $k = k_0, ..., \hat{K} - 1$ while, by construction, $\hat{p}_k = p^*_k$ for all $k \leq k_0 - 1$. We claim that distribution p^* stochastically dominates distribution $\hat{p}$. To see this, fix any $\kappa > \hat{K}$. Clearly, $\sum_{k=\kappa}^{K} \hat{p}_k = 0 \leq \sum_{k=\kappa}^{K} p^*_k$. Now, fix $\kappa \leq \hat{K}$.

$$\sum_{k=\kappa}^{K} \hat{p}_k = 1 - \sum_{k=0}^{\kappa-1} \hat{p}_k \leq 1 - \sum_{k=0}^{\kappa-1} p^*_k = \sum_{k=\kappa}^{K} p^*_k \tag{A.4}$$

where the inequality uses the fact that $\hat{p}_k \geq p^*_k$ for all $k = 0, ..., \kappa - 1$. Importantly, the above inequality is strict for all $\kappa \in \{k_0 + 1, ..., \hat{K}\}$ since $\hat{p}_k > p^*_k$ for all $k = k_0, ..., \hat{K} - 1$.[A.37] It is also strict for any $\kappa \geq \hat{K} + 1$ as long as $p^*_\kappa > 0$ since in that case the LHS is simply 0 while the RHS is strictly positive. In particular, given our assumption that $p^*_k > 0$ for some $k > k^{**}$, it must be that $p^*_{k^{**}+1} > 0$ and so we have

$$\sum_{k=\kappa}^{K} \hat{p}_k < \sum_{k=\kappa}^{K} p^*_k \tag{A.5}$$

for $\kappa = \max\{k_0 + 1, k^{**} + 1\}$.

Now, we show that the value of the objective in $[\mathcal{L}_{\xi^*}]$ strictly increases when we replace solution p^* by $\hat{p}$. We have to show that

$$\sum_{k=0}^{K} \hat{p}_k f(k;\xi^*) > \sum_{k=0}^{K} p^*_k f(k;\xi^*).$$

[A.37] Recall that, by construction, $k_0 + 1 \leq \hat{K}$.

Since $\hat{p}_k = p_k^*$ for all $k \leq k_0 - 1$, this is equivalent to showing the following inequality

$$\sum_{k=k_0}^{K} \hat{p}_k f(k; \xi^*) > \sum_{k=k_0}^{K} p_k^* f(k; \xi^*) \tag{A.6}$$

Now, define a function $\varphi : \mathbb{N} \to \mathbb{R}$ as follows. On the set $\{0, ..., k_0 - 1\}$, function φ equals to an arbitrary decreasing function which is greater than $f(k_0; \xi^*)$. In addition, for all $k \geq k_0$, $\varphi(k) = f(k; \xi^*)$. Since $k_0 > k^*$, by Lemma A.5, this function is weakly decreasing and it is strictly decreasing from k to $k+1$ for any $k \geq \max\{k_0, k^{**}\}$. Thus, $\varphi(\kappa - 1) > \varphi(\kappa)$ for $\kappa = \max\{k_0 + 1, k^{**} + 1\}$. Now, we know that p^* stochastically dominates $\hat{p}$, that inequality (A.5) holds at $\kappa = \max\{k_0 + 1, k^{**} + 1\}$. and that $\varphi(\kappa - 1) > \varphi(\kappa)$. Applying Lemma A.7, the expectation of φ with respect to $\hat{p}$ must be strictly greater than its expectation with respect to p^*, i.e.,

$$\sum_{k=0}^{K} (\hat{p}_k - p_k^*)\, \varphi(k) > 0.$$

Since $\hat{p}_k = p_k^*$ for all $k \leq k_0 - 1$, this is equivalent to Equation (A.6). To conclude, $\mathcal{L}(\hat{p}, \xi^*) > \mathcal{L}(p^*, \xi^*)$ which contradicts the fact that (p^*, ξ^*) is a saddle point of $\mathcal{L}(\cdot, \cdot)$.

Case 2 : $p_k^* = 0$ for all $k > k^{**}$. Recall our assumption that there is a pair of states k and k' in the support of p^* satisfying $f(k; \xi^*) \neq f(k'; \xi^*)$. Hence, because $f(\cdot; \xi^*)$ is single-peaked, f must be weakly increasing on the support of p^* and strictly increasing from k to $k+1$ for all $k < k^*$. In particular, this holds at $k = 0$, and so we have $f(0; \xi^*) < f(1; \xi^*)$ and $p_0^* > 0$. Recall that k_0 is the first state in $\{k^* + 1, ..., k^{**} - 1\}$ at which $\mu_k p_k^* < \lambda_{k-1} p_{k-1}^*$ and $p_{k+1}^* > 0$. Build $\hat{p}$ as follows. $\hat{p}_k = p_k^*/Z_1$ for all $k \leq k_0 - 1$ where $Z_1 > 1$, $\hat{p}_k = p_k^*$ for each $k \geq k_0 + 1$ and $\hat{p}_{k_0} = p_{k_0}^* + Z_2$ where $Z_2 \triangleq \sum_{k=0}^{k_0} (p_k^* - \hat{p}_k)$ so that $\hat{p}$ sums up to 1. We pick Z_1 small enough so that $\hat{p}_{k_0}$ remains between 0 and 1 for each k. We show that, for $Z_1 > 1$ small enough, for each $k \leq K$, $\mu_k \hat{p}_k \leq \lambda_{k-1} \hat{p}_{k-1}$. To see that this is true, first fix $k \leq k_0 - 1$ and note that

$$\mu_k \hat{p}_k = \mu_k p_k^*/Z_1 \leq \lambda_{k-1} p_{k-1}^*/Z_1 = \lambda_{k-1} \hat{p}_{k-1}$$

where the inequality simply comes from the fact that p^* is a solution of $[P'_K]$. Now, for $k = k_0$, we have

$$\mu_{k_0} \hat{p}_{k_0} = \mu_{k_0} \left(p_{k_0}^* + Z_2\right) \leq \lambda_{k_0 - 1} p_{k_0 - 1}^*/Z_1 = \lambda_{k_0 - 1} \hat{p}_{k_0 - 1}$$

where the inequality holds for Z_1 small enough since, by assumption, $\mu_{k_0} p_{k_0}^* < \lambda_{k_0 - 1} p_{k_0 - 1}^*$ (and Z_2 vanishes as Z_1 goes to 1). Now, for $k = k_0 + 1$, we have

$$\mu_{k_0+1} \hat{p}_{k_0+1} = \mu_{k_0+1} p_{k_0+1}^* \leq \lambda_{k_0} p_{k_0}^* \leq \lambda_{k_0} (p_{k_0}^* + Z_2) = \lambda_{k_0} \hat{p}_{k_0}.$$

Finally, by construction, for any $k > k_0 + 1$, $\mu_k \hat{p}_k \leq \lambda_{k-1} \hat{p}_{k-1}$ must hold since p^* and $\hat{p}$ coincide.

Now, we show that the value of the objective in $[\mathcal{L}_{\xi^*}]$ strictly increases when we replace solution p^* by $\hat{p}$. To see this, observe first that $\hat{p}$ must stochastically dominate p^*. Indeed, fix any $\kappa > k_0$. Clearly, since $\hat{p}_k = p^*_k$ for all $k \geq k_0 + 1$, $\sum_{k=\kappa}^{K} \hat{p}_k = \sum_{k=\kappa}^{K} p^*_k$. Now, fix $\kappa \leq k_0$.

$$\sum_{k=\kappa}^{K} \hat{p}_k = 1 - \sum_{k=0}^{\kappa-1} \hat{p}_k > 1 - \sum_{k=0}^{\kappa-1} p^*_k = \sum_{k=\kappa}^{K} p^*_k \tag{A.7}$$

where the inequality uses the fact that $\hat{p}_k = p^*_k / Z_1 < p^*_k$ for all $k = 0, ..., \kappa - 1$ (since $Z_1 > 1$ and $p^*_k > 0$ for such k). Now, we have to show that the value of the objective in $[\mathcal{L}_{\xi^*}]$ strictly increases when we replace solution p^*, i.e.,

$$\sum_{k=0}^{K} \hat{p}_k f(k; \xi^*) > \sum_{k=0}^{K} p^*_k f(k; \xi^*).$$

We know that $\hat{p}$ stochastically dominates p^*, that inequality (A.7) holds at $\kappa = 1$ and that $f(0; \xi^*) < f(1; \xi^*)$. In addition, $f(\cdot; \xi^*)$ is nondecreasing on the support of p^* and $\hat{p}$. Hence, this follows from Lemma A.7. ∎

A.2 Infinite dimensional analysis

Let us consider the sequence $\{p^K\}_K$ where for each K, p^K is an optimal solution of problem $[P'_K]$ and exhibits a cutoff policy which is well-defined by Proposition A.6. For each K, we see p^K as a point in $\mathbb{R}^{\mathbb{Z}_+}$ with value 0 on states weakly greater than $K+1$. We will be interested in the limit points of sequence $\{p^K\}_K$. The following statement implies Theorem 1.

Proposition A.7. Sequence $\{p^K\}_K$ has a subsequence which converges to a distribution p^* which is an optimal solution to $[P']$. Further, p^* exhibits a cutoff policy.

We show this result with the following steps. First, we show that the infinite-dimensional problem $[P']$ is well-behaved, i.e., it does have a solution (Proposition A.8). Then, we show that the set of feasible distributions of $[P']$ exhibiting a cutoff-policy is sequentially compact; proving that $\{p^K\}_K$ has a subsequence converging to a point which exhibits a cutoff policy (Proposition A.11). Finally, we argue that any limit point of $\{p^K\}_K$ must be an optimal solution of $[P']$ (Proposition A.12).

A.2.1 Existence of a solution in the infinite-dimensional problem

Our problem $[P']$ can be written as

$$[P'] \qquad \max_{p \in M'} \sum_{k=0}^{\infty} p_k \left[\mu_k((1-\alpha)R + \alpha V) - \alpha C k\right]$$

where $M' \triangleq \{p \in \Delta(\mathbb{Z}_+) : \sum_{k=0}^{\infty} p_k [\mu_k V - Ck] \geq 0, \lambda_k p_k \geq \mu_{k+1} p_{k+1}, \forall k\}$. We prove the following result.

Proposition A.8. The set of optimal solutions of $[P']$ is nonempty.

We start by showing that the objective of the optimization problem is upper semi-continuous (Proposition A.9). In terms of topology, we assume that $\mathbb{Z}_+$ is endowed with the discrete topology and $\Delta(\mathbb{Z}_+)$ with the weak topology. Note that since $\mathbb{Z}_+$ endowed with the discrete topology is a metric space, so is $\Delta(\mathbb{Z}_+)$ by Prokhorov's Theorem. Then, we show that set M' is compact (Proposition A.10). This enough for our purpose. Indeed, by the maximum theorem for upper semi-continuous functions, optimization problem $[P']$ has a solution.

Proposition A.9. The function which maps $p \in \Delta(\mathbb{Z}_+)$ to

$$\sum_{k=0}^{\infty} p_k [\mu_k((1-\alpha)R + \alpha V) - \alpha C k]$$

is upper semi-continuous.

Proof. Consider a sequence $\{p^n\}$ in $\Delta(\mathbb{Z}_+)$ converging to p^*. Since the function $k \mapsto \mu_k((1-\alpha)R + \alpha V) - \alpha Ck$ is continuous (in the discrete topology) and upper bounded[A.38], by Portmanteau's Theorem, $\limsup \sum_{k=0}^{\infty} p_k^n [\mu_k((1-\alpha)R + \alpha V) - \alpha Ck] \leq \sum_{k=0}^{\infty} p_k^* [\mu_k((1-\alpha)R + \alpha V) - \alpha Ck]$ and so we get the upper semi-continuity of our function. ∎

Proposition A.10. Set M' is compact.

Proof. The proof is based on the two lemmas proved below.

Lemma A.8. The set M' is tight.

Proof. We need to show that for any $\varepsilon > 0$, there is n large enough so that any probability measure $p \in M'$ has $\sum_{k=n+1}^{\infty} p_k < \varepsilon$. Proceed by contradiction and assume that there is $\varepsilon > 0$ and a sequence $\{p^n\}_n$ in M' (which satisfies $\sum_{k=0}^{\infty} p_k^n [\mu_k V - Ck] \geq 0$) such that $\sum_{k=n+1}^{\infty} p_k^n > \varepsilon$ for all n. This implies

$$\begin{aligned}\sum_{k=0}^{\infty} p_k^n(\mu_k V - Ck) &= V \sum_{k=0}^{\infty} p_k^n \mu_k - C \sum_{k=0}^{\infty} p_k^n k \\ &\leq V - C \sum_{k=n+1}^{\infty} p_k^n k\end{aligned}$$

[A.38] Recall our assumption that μ_k is uniformly bounded.

$$
\begin{aligned}
&\leq \; V - C(n+1) \sum_{k=n+1}^{\infty} p_k^n \\
&\leq \; V - C(n+1)\varepsilon.
\end{aligned}
$$

Note that for n large enough, the above term must be strictly negative. This contradicts the fact that $\sum_{k=0}^{\infty} p_k^n(\mu_k V - Ck) \geq 0$ for all n. ∎

Lemma A.9. The set M' is closed.

Proof. Recall that since $\Delta(\mathbb{Z}_+)$ is a metric space, $p \in \Delta(\mathbb{Z}_+)$ is a limit point of M if and only if there is a sequence of points in $M \backslash \{p\}$ converging to p. We need to show that any limit point of M' is contained in M'. Take any sequence $\{p^n\}_n$ in M' converging to p^*. We need to show that (1) $\sum_{k=0}^{\infty} p_k^*(\mu_k V - Ck) \geq 0$ and (2) for all k, $\lambda_k p_k^* \geq \mu_{k+1} p_{k+1}^*$.

(1) $\sum_{k=0}^{\infty} p_k^*(\mu_k V - Ck) \geq 0$. Proceed by contradiction and assume that $\sum_{k=0}^{\infty} p_k^*(\mu_k V - Ck) < 0$. By Portmanteau's Theorem, since the function $k \mapsto \mu_k V - Ck$ is bounded above (and trivially continuous in the discrete topology), we must have that $\limsup \sum_{k=0}^{\infty} p_k^n(\mu_k V - Ck) \leq \sum_{k=0}^{\infty} p_k^*(\mu_k V - Ck)$. Hence, since, by assumption, $\sum_{k=0}^{\infty} p_k^*(\mu_k V - Ck) < 0$, it must be that for n large enough, $\sum_{k=0}^{\infty} p_k^n(\mu_k V - Ck) < 0$, a contradiction with the fact that $p^n \in M'$.

(2) For all k, $\lambda_k p_k^* \geq \mu_{k+1} p_{k+1}^*$. By contradiction, assume that for some k, $\lambda_k p_k^* < \mu_{k+1} p_{k+1}^*$. Since p_k^n and p_{k+1}^n converge pointwise to p_k^* and p_{k+1}^*, for n large enough we have $\lambda_k p_k^n < \mu_{k+1} p_{k+1}^n$ which contradicts the fact that p^n is in M'. ∎

Hence, M' is closed and tight. By Prokhorov Theorem, M' must be sequentially compact. Since $\Delta(\mathbb{Z}_+)$ is a metric space, this implies that M' is compact, as claimed. ∎

A.2.2 Completion of the proof of Proposition A.7

Let M'' be the set of p in M' which exhibits a cutoff policy, i.e., such that for some $\hat{K}$, $\lambda_k p_k = \mu_{k+1} p_{k+1}, \forall k = 0, ..\hat{K} - 1$ and $p_k = 0$ for all $k \geq \hat{K} + 1$. Recall that the sequence $\{p^K\}_K$ is defined as for each K, p^K is an optimal solution of $[P'_K]$ and exhibits a cutoff policy. In addition, we recall that for each K, we see p^K as a point in $\mathbb{R}^{\mathbb{Z}_+}$ with value 0 on states weakly greater than $K + 1$. Clearly $\{p^K\}_K$ is a sequence in M''. In the next proposition we show that M'' is (sequentially) compact. This will show that $\{p^K\}_K$ must have a subsequence converging to a point exhibiting a cutoff policy.

Proposition A.11. $\{p^K\}_K$ must have a subsequence converging to a feasible point p^* of $[P']$ exhibiting a cutoff policy. In addition, $p_k^* > 0$ for each $k \leq \min \arg\max_k \mu_k V - Ck$.

Proof. In order to prove the first part of the statement, it is enough to show that M'' is (sequentially) compact. Since M'' is a subset of M' which is compact (Proposition A.10), we only need to show that M'' is closed. Consider a sequence $\{p^n\}$ in M'' converging to p^*.

We need to show that $p^* \in M''$. Since M' is (sequentially) compact, we already know that $p^* \in M'$. Letting $\hat{K}$ be the largest state in the support of p^* (which is potentially $+\infty$ if the support is unbounded), we proceed by contradiction and assume that there exists $k_0 < \hat{K}$ such that $\lambda_{k_0-1} p^*_{k_0-1} > \mu_{k_0} p^*_{k_0}$. Now, simply pick n large enough so that (1) $p^n_k > 0$ for all $k = 0, ..., k_0 + 1$ and (2) $\lambda_{k_0-1} p^n_{k_0-1} > \mu_{k_0} p^n_{k_0}$. This contradicts the assumption that p^n is in M''.

Now, we show the second part of the statement. We just proved that $\{p^K\}_K$ must have a subsequence converging to a feasible point p^* of $[P']$. We show that p^* satisfies $p^*_k > 0$ for each $k \leq \min \arg\max_k \mu_k V - Ck$. First, we simply observe that for any $\xi \geq 0$, $\min \arg\max_k \mu_k V - Ck \leq \min \arg\max f(k;\xi)$. Now, we proceed by contradiction and assume that there is $k_0 \leq \min \arg\max_k \mu_k V - Ck$ such that $p^*_{k_0} = 0$. Let us assume that k_0 is the smallest state satisfying this property, so, in particular, $p^*_{k_0-1} > 0$. This implies that $p^*_{k_0}\mu_{k_0} < p^*_{k_0-1}\lambda_{k_0-1}$. Since $\{p^K\}_K$ converges to p^*, we must have that for K large enough, $p^K_{k_0}\mu_{k_0} < p^K_{k_0-1}\lambda_{k_0-1}$. Note that this contradicts Lemma A.6 since $k_0 \leq \min \arg\max f(k;\xi^*_K)$ and so using single-peakedness of $f(\cdot;\xi^*_K)$, we must have $f(k_0-1;\xi^*_K) < f(k_0;\xi^*_K)$ (where we use the notation (p^K, ξ^*_K) for the saddle point of the Lagrangian in $[P'_K]$). ∎

Finally, we complete the proof of Proposition A.7 via the following proposition.

Proposition A.12. Take any subsequence of $\{p^K\}_K$ converging to a limit p^*. p^* must be an optimal solution of $[P']$.

Proof. In the sequel, we let p^* be the limit of an arbitrary converging subsequence $\{p^K\}_K$. We proceed by contradiction and assume that p^* is not a solution to the infinite dimensional problem. By Proposition A.8, we know that there is a solution to this problem. Let us call it $\bar{p}$. By assumption,

$$\sum_{k=0}^{\infty} \bar{p}_k \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right] > \sum_{k=0}^{\infty} p^*_k \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right]. \tag{A.8}$$

Now, let us note by $\bar{p}^K$ the distribution $\bar{p}$ conditional on $\{0, ..., K\}$, i.e., $\bar{p}^K_k = 0$ for all $k \geq K+1$ while $\bar{p}^K_k = \bar{p}_k \Big/ \sum_{k=0}^{K} \bar{p}_k$ for all $k \leq K$. We claim that

$$\lim \sum_{k=0}^{\infty} \bar{p}^K_k \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right] = \sum_{k=0}^{\infty} \bar{p}_k \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right].$$

Indeed, by construction, for each K,

$$\sum_{k=0}^{\infty} \bar{p}^K_k \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right] = \sum_{k=0}^{K} \bar{p}_k \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right] \Bigg/ \sum_{k=0}^{K} \bar{p}_k \,.$$

Taking limits on both sides as K increases gives us (and using the fact that $\lim_{K\to\infty}\sum_{k=0}^{K}\bar{p}_k = 1$),

$$\lim \sum_{k=0}^{\infty} \bar{p}_k^K \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right] = \sum_{k=0}^{\infty} \bar{p}_k \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right],$$

as claimed.

Now, using Equation (A.8), for K large enough, we must have

$$\sum_{k=0}^{\infty} \bar{p}_k^K \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right] > \sum_{k=0}^{\infty} p_k^* \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right] + \varepsilon \tag{A.9}$$

for some $\varepsilon > 0$. Now, since $\{p^K\}_K$ converges weakly to p^*, by Proposition A.9,

$$\limsup \sum_{k=0}^{\infty} p_k^K \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right] \leq \sum_{k=0}^{\infty} p_k^* \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right].$$

Hence, we must have that for K large enough,

$$\sum_{k=0}^{\infty} p_k^* \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right] + \varepsilon > \sum_{k=0}^{\infty} p_k^K \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right]. \tag{A.10}$$

Using Equation (A.9) and (A.10), we conclude that for K large enough,

$$\sum_{k=0}^{\infty} \bar{p}_k^K \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right] > \sum_{k=0}^{\infty} p_k^K \left[\mu_k \left((1-\alpha) R + \alpha V\right) - \alpha C k\right].$$

This contradicts that the fact that p^K is an optimal solution of $[P_K']$ since $\bar{p}^K$ is feasible in this problem. ∎

B Proof of Proposition 1

In the sequel, given $K \in \mathbb{Z}_+ \cup \{\infty\}$ and $x \in (0,1]$, we consider the probability distribution over $\mathbb{Z}_+$ defined by

$$p_k = p_0 \prod_{\ell=1}^{k} \frac{\lambda_{\ell-1}}{\mu_\ell} \tag{B.11}$$

for any $k = 1, ..., K-1$ and if $K < \infty$

$$p_K = p_0 \prod_{\ell=1}^{K-1} \frac{\lambda_{\ell-1}}{\mu_\ell} \frac{\lambda_{K-1} x}{\mu_K}$$

while p_0 is defined to ensure that the total mass of p_k's is equal to 1. Given $K \in \mathbb{Z}_+ \cup \{\infty\}$ and x, whenever well-defined, such a distribution will be denoted by $p(K, x)$.[B.39] By Theorem 1 we know that, when μ is regular, there is an optimal solution p^* of $[P']$ which can be implemented by a cutoff policy. Recall our observation that this implies that (B′) binds for all $k = 0, ..., K-1$ and holds with weak inequality for $k = K-1$ where K is the largest state in the support of p^*. A simple inductive argument yields that p^* is equal to $p(K^*, x)$ for some $K \in \mathbb{Z}_+ \cup \{\infty\}$ and $x^* \in (0, 1]$. We say $p(K^*, x^*)$ is optimal in that case.

Central to our analysis is the following function $\psi : \mathbb{Z}_+ \cup \{\infty\} \to \mathbb{R}$ defined as

$$\psi(K) \triangleq \sum_{k=1}^{K} \prod_{\ell=1}^{k} \left(\frac{\lambda_{\ell-1}}{\mu_\ell} \right) [\mu_k V - Ck] .$$

Proposition 1 follows from the two propositions below. Recall that the environment is **generic** if $\psi(K) \neq 0$ for all $K \in \mathbb{Z}_+ \cup \{\infty\}$. Our first proposition is stated as follows.

Proposition B.13. Assume μ is regular. Suppose $p(K^*, x^*)$ is optimal. If $\alpha = 0$, then $K^* < \infty$ if and only if $\psi(\infty) < 0$. If $\psi(\infty) < 0$, then $K^* < \infty$ and (IR) binds if $\alpha > 0$ is sufficiently small or if $\alpha < 1$ and R is sufficiently large. If (IR) binds and the environment is generic, then $x^*_{K^*-1} \in (0, 1)$.

Our second proposition states that the maximum queue length at the optimal policy is finite whenever $\alpha > 0$.

Proposition B.14. Assume μ is regular and $\alpha \in (0, 1]$. Let $p(K^*, x^*)$ be the optimal policy. We have $K^* < \infty$.

B.1 Proof of Proposition B.13

We first prove that, as K and x increase, $p(K, x)$ places more mass to higher states.

Lemma B.10. The following must hold.

1. Fix $K' > K$ and any $x, x' \in (0, 1]$. Whenever well-defined, $p(K', x')$ dominates $p(K, x)$ in the likelihood ratio order and the mass on state 0 is strictly smaller under $p(K', x')$ than under $p(K, x)$.
2. Assume $K < \infty$. Fix any $x, x' \in (0, 1]$ with $x' > x$. We have that $p(K, x')$ dominates $p(K, x)$ in the likelihood ratio order and the mass on state 0 is strictly smaller under $p(K, x')$ than under $p(K, x)$.

[B.39] We note that for $K = \infty$, p_0 may not always be well-defined.

Proof. We start by showing Part 1. Fix any $K' > K$ and any $x, x' \in (0,1]$. To ease notations, we let p' and p be $p(K', x')$ and $p(K, x)$ respectively. By construction, we have for all $k = 1, ..., K-1$

$$\frac{p'_k}{p'_{k-1}} = \frac{\lambda_{k-1}}{\mu_k} = \frac{p_k}{p_{k-1}}$$

while

$$\frac{p'_K}{p'_{K-1}} = \frac{\lambda_{K-1}}{\mu_K} \geq \frac{\lambda_{K-1}x}{\mu_K} = \frac{p_K}{p_{K-1}}.$$

Finally, for any $k \geq K+1$,

$$p'_k p_{k-1} \geq p'_{k-1} p_k = 0.$$

Thus, p' dominates p in the likelihood ratio order. Now, we argue that $p'_0 < p_0$. By construction, we have

$$1 = p_0 + \sum_{k=1}^{K} p_k \leq p_0 + p_0 \sum_{k=1}^{K} \prod_{\ell=1}^{k} \left(\frac{\lambda_{\ell-1}}{\mu_\ell} \right) \tag{B.12}$$

and

$$1 = p'_0 + \sum_{k=1}^{K'} p'_k > p'_0 + p'_0 \sum_{k=1}^{K'-1} \prod_{\ell=1}^{k} \left(\frac{\lambda_{\ell-1}}{\mu_\ell} \right) \geq p'_0 + p'_0 \sum_{k=1}^{K} \prod_{\ell=1}^{k} \left(\frac{\lambda_{\ell-1}}{\mu_\ell} \right) \tag{B.13}$$

where the strict inequality follows since $p'_{K'} > 0$. Equations (B.12) and (B.13) together imply that $p'_0 < p_0$, as claimed.

We next prove Part 2. Assume $K < \infty$. Fix any $x, x' \in (0,1]$ with $x' > x$. Again to ease notations, we let p' and p be $p(K, x')$ and $p(K, x)$ respectively. Again, by construction, we have for all $k = 1, ..., K-1$

$$\frac{p'_k}{p'_{k-1}} = \frac{\lambda_{k-1}}{\mu_k} = \frac{p_k}{p_{k-1}}$$

while

$$\frac{p'_K}{p'_{K-1}} = \frac{\lambda_{K-1}x'}{\mu_K} \geq \frac{\lambda_{K-1}x}{\mu_K} = \frac{p_K}{p_{K-1}}.$$

We conclude that p' dominates in the likelihood ratio order distribution p. Now, we argue that $p'_0 < p_0$. By construction, we have

$$1 = p_0 + p_0 \sum_{k=1}^{K-1} \prod_{\ell=1}^{k} \left(\frac{\lambda_{\ell-1}}{\mu_\ell} \right) + p_0 \prod_{\ell=1}^{K-1} \left(\frac{\lambda_{\ell-1}}{\mu_\ell} \right) \left(\frac{\lambda_{K-1}x}{\mu_K} \right) \tag{B.14}$$

and

$$1 \quad = \quad p'_0 + p'_0 \sum_{k=1}^{K-1} \prod_{\ell=1}^{k} \left(\frac{\lambda_{\ell-1}}{\mu_\ell} \right) + p'_0 \prod_{\ell=1}^{K-1} \left(\frac{\lambda_{\ell-1}}{\mu_\ell} \right) \left(\frac{\lambda_{K-1}x'}{\mu_K} \right)$$

$$> \quad p_0' + p_0' \sum_{k=1}^{K-1} \prod_{\ell=1}^{k} \left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right) + p_0' \prod_{\ell=1}^{K-1} \left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right) \left(\frac{\lambda_{K-1} x}{\mu_K}\right) \tag{B.15}$$

where the strict inequality comes from the fact that $x' > x$. Equations (B.14) and appendix B.1 together imply that $p_0' < p_0$, as claimed. ∎

We show that, given the regularity of μ, function ψ is single-peaked: $\psi(K-1) < \psi(K)$ for all $K \leq \min \arg\max_{K \in \mathbb{Z}_+} \psi(K)$ while $\psi(K+1) < \psi(K)$ for all $K \geq \max \arg\max_{K \in \mathbb{Z}_+} \psi(K)$. We recall our convention that if $\arg\max_{K \in \mathbb{Z}_+} \psi(K)$ is empty, then $\min \arg\max_{K \in \mathbb{Z}_+} \psi(K)$ is set to $+\infty$.

Lemma B.11. Assume μ is regular. Function ψ is single-peaked.

Proof. Single-peakedness of ψ is equivalent to: $\psi(K-1) \leq (<)\psi(K)$ implies $\psi(K'-1) \leq (<)\psi(K')$ for all $K' \leq K$. We just prove the strict inequality version; the argument for the weak inequality version is identical. Assume that $\psi(K-1) < \psi(K)$. This is equivalent to

$$\prod_{\ell=1}^{K} \left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right) [\mu_K V - CK] > 0$$

which in turn is equivalent to $\mu_K V - CK > 0$. We claim that this implies $\mu_{K'} V - CK' > 0$ for any $K' \leq K$—which by the above reasoning will imply $\psi(K'-1) < \psi(K')$.

First, by the regularity of μ, recall Lemma A.5 (with $\alpha = 1$ and $\xi = 0$) which proves that $h(K) \triangleq \mu_K V - CK$ is single-peaked. Now, observe that function h is equal to 0 at $K = 0$. Hence, if $h(K) = \mu_K V - CK > 0$, given single-peakedness of h, we must have $h(K') = \mu_{K'} V - CK' > 0$ for any $K' = 1, ..., K$, as claimed. ∎

We define $\bar{K} \triangleq \inf\{K' \in \mathbb{Z}_+ : \psi(K') < 0\}$ and $\bar{K} = \infty$ if $\psi(\infty) \geq 0$. Given Lemma B.11 (together with $\psi(0) = 0$), we have that $\psi(K) \geq 0$ if and only if $K < \bar{K}$ whenever $\bar{K} < \infty$. In the sequel, given $K \in \mathbb{Z}_+ \cup \{\infty\}$, $x \in (0, 1]$ and $\alpha \in [0, 1]$, whenever $p(K, x)$ is well-defined, we let $\Phi(K, x; \alpha) \triangleq \sum_{k=1}^{\infty} p_k(K, x) \left[(1-\alpha)\mu_k R + \alpha(\mu_k V - Ck)\right]$ be the objective in $[P']$. Note that K can be infinite so that $p(K, x)$ may not always be well-defined. In the online appendix Section S.2, we show that $p(\infty, x)$ is well-defined whenever $\bar{K} = \infty$ and so for any x and α, $\Phi(\bar{K}, x; \alpha)$ is well-defined for any value of $\bar{K}$.

In Proposition B.13 below we will show that (for α small enough) the optimal policy is $p(K^*, x^*)$ for $K^* \in \{\bar{K} - 1, \bar{K}\}$ and some $x^* \in (0, 1]$. In order to show this, we will use the following straightforward corollary of Lemma B.10.

Corollary B.4. Assume μ is regular. Fix any $\hat{K} \in \{\bar{K} - 1, \bar{K}\}$ and $\hat{x} \in (0, 1]$. We have

1. For $K < \hat{K}$ and $x \in (0, 1]$

$$\Phi(\hat{K}, \hat{x}; 0) > \Phi(K, x; 0).$$

2. For $x, x' \in (0,1]$ with $x' > x$

$$\Phi(\hat{K}, x'; 0) > \Phi(\hat{K}, x; 0).$$

Proof. Given Lemma B.10 and the fact that, by regularity, μ_k is nondecreasing in k, it is straightforward to show that each of the two inequalities must hold weakly. The fact that it holds with strict inequality comes from Lemma A.7 together with the observation that $\mu_1 R > \mu_0 R = 0$. ∎

Corollary B.4 above is proved for $\alpha = 0$. We show below that, provided that $\bar{K} < \infty$, a similar result applies when α is small enough.

Lemma B.12. Assume μ is regular. Let $\bar{K} < \infty$. Fix any $\hat{K} \in \{\bar{K}-1, \bar{K}\}$ and $\hat{x} \in (0,1]$. There is $\hat{\alpha} > 0$ such that for any $\alpha < \hat{\alpha}$, we have

1. For $K < \hat{K}$ and $x \in (0,1]$

$$\Phi(\hat{K}, \hat{x}; \alpha) > \Phi(K, x; \alpha).$$

2. For $x, x' \in (0,1]$ with $x' > x$

$$\Phi(\hat{K}, x'; \alpha) > \Phi(\hat{K}, x; \alpha).$$

Proof. We first make the simple observation that $\Phi(K, x; \alpha) \geq 0$ for any x and $K < \bar{K}$. In order to show this, observe that for any x and $K < \bar{K}$, $\Phi(K, x; \alpha)$ is equal to

$$p_0(K,1) \sum_{k=1}^{K-1} \prod_{\ell=1}^{k} \left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right) g(k,\alpha) + p_0(K,1) \prod_{\ell=1}^{K-1} \left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right) \left(\frac{\lambda_{K-1} x}{\mu_K}\right) g(K,\alpha)$$

where $g(k,\alpha) \triangleq (1-\alpha)\mu_k R + \alpha(\mu_k V - Ck)$. Given that for all k, $\mu_k R \geq 0$, the above expression is greater than

$$\alpha p_0(K,1) \left[\sum_{k=1}^{K-1} \prod_{\ell=1}^{k} \left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right) (\mu_k V - Ck) + \prod_{\ell=1}^{K-1} \left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right) \left(\frac{\lambda_{K-1} x}{\mu_K}\right) (\mu_K V - CK) \right].$$

Now, if $\mu_K V - CK \leq 0$ then the above term is weakly greater than $\alpha p_0(K,1)\psi(K)$ which is weakly greater than 0 since $K < \bar{K}$. If, on the contrary, $\mu_K V - CK \geq 0$ then the above term is weakly greater than $\alpha p_0(K,1)\psi(K-1)$ which for the same reason must be no less than 0. We conclude that $\Phi(K, x; \alpha) \geq 0$ for any x and $K < \bar{K}$.

Now, let us fix $\hat{K} \in \{\bar{K}-1, \bar{K}\}$ and any $\hat{x} \in (0,1]$. By Corollary B.4, we know that $\Phi(\hat{K}, \hat{x}; 0) > \Phi(K, 1; 0) \geq \Phi(K, x; 0)$ for any $K < \hat{K}$ and $x \in (0,1]$. Consider any $K < \hat{K}$. By continuity, $\lim_{\alpha \to 0} \Phi(\hat{K}, \hat{x}; \alpha) = \Phi(\hat{K}, \hat{x}; 0)$. So for α small enough,

$$\Phi(\hat{K}, \hat{x}; \alpha) > \Phi(K, 1; 0). \tag{B.16}$$

Now, for each $\alpha \in [0,1]$, let $x(\alpha) \in \arg\max_{x\in[0,1]} \Phi(K,x;\alpha)$ which by continuity of $\Phi(K,\cdot;\alpha)$ and compactness of $[0,1]$ is well-defined. By Berge's Maximum Theorem and Corollary B.4, $\lim_{\alpha\to 0} \Phi(K,x(\alpha);\alpha) = \Phi(K,1;0)$. This, together with Equation (B.16), imply that there is $\hat{\alpha} > 0$ such that for any $\alpha < \hat{\alpha}$, $\Phi(\hat{K},\hat{x};\alpha) > \Phi(K,x(\alpha);\alpha) \geq \Phi(K,x;\alpha)$ for any $x \in [0,1]$. Since there are finitely many positive integers $K < \hat{K}$, we obtain that there is $\hat{\alpha} > 0$ such that for any $\alpha < \hat{\alpha}$,

$$\Phi(\hat{K},\hat{x};\alpha) > \Phi(K,x;\alpha) \tag{B.17}$$

for any $x \in [0,1]$ and any $K < \hat{K}$.

Now, for part 2., we need to show that for α small enough and for $x' > x$

$$\Phi(\hat{K},x';\alpha) > \Phi(\hat{K},x;\alpha).$$

We claim that for a given $\alpha \in [0,1]$, the above inequality holds true for any $x' > x$ if

$$\Phi(\hat{K},1;\alpha) > \Phi(\hat{K}-1,1;\alpha).$$

This would be sufficient for our purpose since by Equation (B.17) (setting $\hat{x} = x = 1$), for α small enough,

$$\Phi(\hat{K},1;\alpha) > \Phi(\hat{K}-1,1;\alpha).$$

Let us note that $\Phi(\hat{K},1;\alpha) > \Phi(\hat{K}-1,1;\alpha)$ means

$$p_0(\hat{K},1)\sum_{k=1}^{\hat{K}}\prod_{\ell=1}^{k}\left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right)g(k,\alpha) > p_0(\hat{K}-1,1)\sum_{k=1}^{\hat{K}-1}\prod_{\ell=1}^{k}\left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right)g(k,\alpha)$$

where we recall that $g(k,\alpha) = (1-\alpha)\mu_k R + \alpha(\mu_k V - Ck)$. Since, by Lemma B.10, $p_0(\hat{K},1) < p_0(\hat{K}-1,1)$ and since, as we already argued, $\Phi(\hat{K}-1,1;\alpha) \geq 0$ because $\hat{K} \leq \bar{K}$, the above expression implies

$$p_0(\hat{K},1)\sum_{k=1}^{\hat{K}}\prod_{\ell=1}^{k}\left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right)g(k,\alpha) > p_0(\hat{K},1)\sum_{k=1}^{\hat{K}-1}\prod_{\ell=1}^{k}\left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right)g(k,\alpha).$$

This implies $g(\hat{K},\alpha) > 0$. Now, simple algebra yields that $\Phi(\hat{K},x';\alpha) > \Phi(\hat{K},x;\alpha)$ is equivalent to

$$\prod_{\ell=1}^{\hat{K}-1}\left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right)\left(\frac{\lambda_{K-1}x'}{\mu_K}\right)g(\hat{K},\alpha) > \prod_{\ell=1}^{\hat{K}-1}\left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right)\left(\frac{\lambda_{K-1}x}{\mu_K}\right)g(\hat{K},\alpha).$$

The above expression holds true whenever $x' > x$ and $\Phi(\hat{K},1;\alpha) > \Phi(\hat{K}-1,1;\alpha)$ since, as we just showed, the latter implies $g(\hat{K},\alpha) > 0$. ∎

Completion of the proof of Proposition B.13. Fix $\alpha = 0$. Assume that $\psi(\infty) \geq 0$. This means that (IR) is satisfied at $\bar{K} = \infty$. Corollary B.4 yields that $K^* = \infty$.

Now, assume that $\psi(\infty) < 0$ and note that this implies that $\bar{K} < \infty$. We consider two cases.

<u>Case 1.</u> $\psi(\bar{K}-1) = 0$. Pick $\hat{\alpha}$ as defined in Lemma B.12 where $\hat{K} = \bar{K} - 1$ and $\hat{x} = 1$. we show that if $\alpha < \hat{\alpha}$ then $K^* < \infty$ and (IR) binds. We claim that the optimal policy has $K^* = \bar{K} - 1$ and $x^* = 1$, i.e., it is equal to $p(\bar{K}-1, 1)$. Indeed, for any $x \in (0, 1]$ and $K \geq \bar{K}$,

$$\sum_{k=1}^{K} p_k(K,x)\,[\mu_k V - Ck] < p_0(K,x) \sum_{k=1}^{\bar{K}-1} \prod_{\ell=1}^{k} \left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right) [\mu_k V - Ck] = \psi(\bar{K}-1) = 0$$

where the strict inequality holds since $\psi(\bar{K}) < 0$. The inequality implies $\mu_K V - CK < 0$ for all $K \geq \bar{K}$.[B.40] Hence, (IR) is violated at $p(K, x)$. Now, because $\psi(\bar{K}-1) = 0$ implies that (IR) holds for $p(\bar{K}-1, 1)$, Lemma B.12 completes the argument.

<u>Case 2.</u> $\psi(\bar{K}-1) > 0$. Note that $\sum_{k=1}^{\bar{K}} p_k(\bar{K}, x)\,[\mu_k V - Ck]$ tends to $\psi(\bar{K}-1) > 0$ when x goes to 0 and to $\psi(\bar{K}) < 0$ when x goes to 1 and it is nonincreasing in x. Thus, there is a unique $\bar{x} \in (0, 1)$ such that

$$\sum_{k=1}^{\bar{K}-1} \prod_{\ell=1}^{k} \left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right) [\mu_k V - Ck] + \prod_{\ell=1}^{\bar{K}-1} \left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right) \left(\frac{\lambda_{\bar{K}-1} x}{\mu_{\bar{K}}}\right) \left[\mu_{\bar{K}} V - C\bar{K}\right] = 0. \tag{B.18}$$

Pick $\hat{\alpha}$ as defined in Lemma B.12 where $\hat{K} = \bar{K}$ and $\hat{x} = \bar{x}$. We show that if $\alpha < \hat{\alpha}$ then $K^* < \infty$ and (IR) binds. We claim that the optimal policy has $K^* = \bar{K}$ and $x^* = \bar{x}$, i.e., it is equal $p(\bar{K}, \bar{x})$. Our choice of $\bar{x} \in (0, 1)$ ensures that (IR) is binding. Using a similar argument as in Case 1, it is easy to show that for any $x > \bar{x}$, (IR) is violated at $p(\bar{K}, x)$ and similarly at $p(K, x)$ for $K > \bar{K}$ and $x \in (0, 1]$. We conclude using Lemma B.12.

Given the above arguments in Case 1. and 2., we conclude that when $\psi(\infty) < 0$, for any $\alpha > 0$ small enough, $K^* < \infty$ and (IR) binds. Since R is interchangeable with $1/\alpha$ in the objective of $[P']$, as long as $\alpha < 1$, the same conclusion holds for any $\alpha < 1$, if $\psi(\infty) < 0$ and R is large enough.[B.41]

[B.40] Indeed, since ψ is single-peaked, $\psi(0) = 0$ and $\psi(\bar{K}) < 0$, it must be that ψ is strictly decreasing over $\{\bar{K}, ...\}$. Simple algebra shows that this implies $\mu_K V - CK < 0$ for any $K \geq \bar{K}$.

[B.41] To see this, assume $\psi(\infty) < 0$ and fix any $\alpha < 1$, and let the service provider's profit be $R = \beta\hat{R}$, where $\hat{R} > 0$ is an arbitrary positive number. The objective function of $[P']$ can be rewritten as

$$\frac{(1-\alpha)\beta}{(1-\alpha)\beta + \alpha} \hat{R}\mu_k + \frac{\alpha}{(1-\alpha)\beta + \alpha} (V\mu_k - Ck).$$

Clearly, $[P']$ with R is the same as $[P']$ with revenue $\hat{R}$ and weight (on agents' welfare) $\alpha/((1-\alpha)\beta + \alpha)$. Since $\alpha < 1$, this latter weight goes to 0 as $\beta \to \infty$, or as $R \to \infty$. Hence, given our previous result, for R is sufficiently large, $K^* < \infty$ and (IR) binds.

Finally, suppose (IR) binds at the optimal policy $p(K^*, x^*)$. If the environment is generic, $\psi(\infty) \neq 0$ and so $K^* \neq \infty$. Given this, if (IR) is binding then $x^* = \bar{x} \in (0, 1]$ as defined in Equation (B.18). Since the environment is generic, $\psi(\bar{K} - 1) \neq 0$ and so $\bar{x} \neq 1$. ∎

B.2 Proof of Proposition B.14

Define function $g : \mathbb{Z}_+ \to \mathbb{R}$ as

$$g(K) \triangleq (1 - \alpha)\, \mu_K R + \alpha \left[\mu_K V - CK\right].$$

Lemma A.5 (with $\xi = 0$) implies that function g is single-peaked when μ is regular. Define $\bar{K}_2 \triangleq \sup\{K' : g(K') \geq 0\}$, where $\bar{K}_2 \triangleq \infty$ whenever $g(K') \geq 0$ for all $K' \in \mathbb{Z}_+$. Observe that, when μ is regular, because g is single-peaked (and since $g(0) = 0$), $g(K) \geq 0$ if and only if $K \leq \bar{K}_2$. We state the following result.

Lemma B.13. Assume μ is regular. Let $p(K^*, x^*)$ be the optimal policy. We have $K^* \leq \bar{K}_2$.

Proof. The value of the objective of $[P']$ is given by

$$\sum_{k=1}^{K^*} p_k(K^*, x^*) g(k) = p_0(K^*, x^*) \sum_{k=1}^{K^*-1} \prod_{\ell=1}^{k} \frac{\lambda_{\ell-1}}{\mu_\ell} g(k) + p_0(K^*, x^*) \prod_{\ell=1}^{K^*-1} \frac{\lambda_{\ell-1}}{\mu_\ell} \frac{\lambda_{K^*-1} x^*}{\mu_{K^*}} g(K^*).$$

By way of contradiction, assume that $K^* > \bar{K}_2$. As we already stated, by definition of $\bar{K}_2$ and single-peakedness of g, it must be that $g(K') < 0$ for $K' = \bar{K}_2 + 1, ..., K^*$. Consider the distribution $p(\bar{K}_2, 1)$. Compared to $p(K^*, x^*)$, this distribution removes all weight on negative values and, for each positive value, increases its weight. This must strictly increase the value of the objective.

Now, it remains to show that (IR) is satisfied under $p(\bar{K}_2, 1)$. Since $p(K^*, x^*)$ is optimal, (IR) holds at $p(K^*, x^*)$:

$$p_0(K^*, x^*) \sum_{k=1}^{K^*-1} \prod_{\ell=1}^{k} \left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right) h(k) + p_0(K^*, x^*) \prod_{\ell=1}^{K^*-1} \left(\frac{\lambda_{\ell-1}}{\mu_\ell}\right) \left(\frac{\lambda_{K^*-1} x^*}{\mu_{K^*}}\right) h(K^*) \geq 0$$

where we recall that $h(k) = \mu_k V - Ck$. From this, it follows that $\psi(K^*) \geq 0$ if $h(K^*) \geq 0$ and that $\psi(K^* - 1) \geq 0$ if $h(K^*) \leq 0$. Since, by Lemma B.11, ψ is single-peaked and $\psi(0) = 0$, this implies that $\psi(K') \geq 0$ for any $K' < K^*$. In particular, $\psi(\bar{K}_2) \geq 0$. This implies that (IR) holds at $p(\bar{K}_2, 1)$. This contradicts our assumption that $p(K^*, x^*)$ is optimal. ∎

Completion of the proof of Proposition B.14. Given Lemma B.13, it is enough to prove that $\bar{K}_2 < \infty$. This follows since $\alpha \in (0, 1]$ and μ_k is uniformly bounded, which imply $g(K) \to -\infty$ a $K \to \infty$. ∎

C Proofs from Section 5: FCFS with No Information

C.1 Proof of Lemma 1

The expected waiting time satisfies the following recursion. The agent in the first position has expected waiting time

$$\tau_1^* = (q_1^* dt)dt + [1 - q_1^* dt](\tau_1^* + dt) + o(dt),$$

since he waits for dt period with probability $q_1^* dt$ and for $\tau_1^* + dt$ periods with the remaining probability. Letting $dt \to 0$, we get

$$\tau_1^* = 1/q_1^* = 1/\mu_1.$$

More generally, the agent in queue position ℓ waits for

$$\tau_\ell^* = (q_\ell^* dt)dt + \left[1 - \sum_{j=1}^{\ell} q_j^* dt\right](\tau_\ell^* + dt) + \left(\sum_{j=1}^{\ell-1} q_j^* dt\right)(\tau_{\ell-1}^* + dt) + o(dt),$$

since he is served in dt period with probability $q_\ell^* dt$, in $\tau_\ell^* + dt$ periods with probability $1 - \sum_{j=1}^{\ell} q_j^* dt$ (when nobody before him is served), and in $\tau_{\ell-1}^* + dt$ periods with probability $\sum_{j=1}^{\ell-1} q_j^* dt$ (when somebody before him is served).

The recursion equations yield a unique solution:

$$\tau_\ell^* = \frac{\ell}{\sum_{j=1}^{\ell} q_j^*} = \frac{\ell}{\mu_\ell},$$

where the last equality follows from feasibility.

Part (ii) of regularity implies that q_ℓ^* is nonincreasing in ℓ. Therefore, for each ℓ

$$\tau_{\ell+1}^* - \tau_\ell^* = \frac{\sum_{j=1}^{\ell} q_j^* - \ell q_{\ell+1}^*}{(\sum_{j=1}^{\ell} q_j^*)(\sum_{j=1}^{\ell+1} q_j^*)} \geq 0.$$

Hence, it follows that τ_ℓ^* is nonincreasing in ℓ. Further, if $2\mu_1 > \mu_2$, then $q_1^* > q_2^* \geq q_\ell^*$ for all $\ell \geq 2$. Then, the above inequality becomes strict for all ℓ, which proves the last statement. ∎

C.2 Proof of Lemma 2

Recall the optimality of the cutoff policy means $x_k^* = 1$ for all $k = 0, ..., K^* - 2$ and $x_k^* = 0$ for all $k > K^* - 1$, and $y_{k,\ell}^* = 0$ for all (k, ℓ). Substitute these into (B). Use the resulting equations to rewrite (1):

$$\tilde{\gamma}_\ell^0 = \frac{p_\ell^* \mu_\ell}{\sum_{i=1}^{K^*} p_i^* \mu_i}, \forall \ell = 1, ..., K^*.$$

An agent's expected payoff when joining the queue after being recommended to do so is:

$$\begin{aligned} V - C\sum_{k=1}^{K^*} \tilde{\gamma}_k^0 \cdot \tau_k^* =& V - C\frac{\sum_{k=1}^{K^*} p_k^* \mu_k \cdot \tau_k^*}{\sum_{i=1}^{K^*} p_i^* \mu_i} \\ =& V - C\frac{\sum_{k=1}^{K^*} p_k^* k}{\sum_{i=1}^{K^*} p_i^* \mu_i} \\ =& \left(\frac{1}{\sum_{i=1}^{K^*} p_i^* \mu_i}\right) \sum_{k=1}^{K^*} p_k^* \left(\mu_k V - kC\right), \end{aligned}$$

where the first equality is from the preceding observation and the second equality follows from Lemma 1. Since $\sum_{i=1}^{K^*} p_i^* \mu_i > 0$, (IC_0) holds if and only if (IR) holds. ∎

C.3 Proof of Lemma 3

In the sequel, we want to show the following lemma which implies Lemma 3.

Lemma C.14. Assume that $\lambda_k - \lambda_{k-1} \leq \mu_k - \mu_{k-1}$ for each k. Consider any (x, y, p) satisfying (B) and (IR) where x_k is nonincreasing in k and $y_{k,\ell} = 0$ for each ℓ, k. Then, for all $\ell \in \{2, ..., K^*\}$, r_ℓ^t is nonincreasing in t for all $t \geq 0$.

In the sequel, we let the effective arrival rate be $\tilde{\lambda}_k \triangleq \lambda_k x_k$ for each k. The following straightforward lemma is stated without proof.

Lemma C.15. Assume that $\lambda_k - \lambda_{k-1} \leq \mu_k - \mu_{k-1}$ for each k. Then, $\tilde{\lambda}_k - \tilde{\lambda}_{k-1} \leq \mu_k - \mu_{k-1}$ holds for each k if x_k is nonincreasing in k.

Consider (x, y, p) satisfying (B) and (IR). Further assume that x_k is nonincreasing in k and $y_{k,\ell} = 0$ for each ℓ, k. We now go through the proof of Lemma C.14. We let $\bar{K}$ be the largest state in the support of p (which can potentially be infinite). We first study the dynamics for the case with $\bar{K} < \infty$. The proof for $\bar{K} = \infty$ requires more care and will be provided thereafter.

C.3.1 The case of $\bar{K} < \infty$.

Using (2), we write for each such $\ell \geq 2$,

$$r_\ell^{t+dt} = \frac{\tilde{\gamma}_\ell^{t+dt}}{\tilde{\gamma}_{\ell-1}^{t+dt}} = \frac{(1-\mu_\ell dt)\tilde{\gamma}_\ell^t + \mu_\ell dt \tilde{\gamma}_{\ell+1}^t}{(1-\mu_{\ell-1}dt)\tilde{\gamma}_{\ell-1}^t + \mu_{\ell-1}dt\tilde{\gamma}_\ell^t} + o(dt) = \frac{1-\mu_\ell dt + \mu_\ell dt r_{\ell+1}^t}{(1-\mu_{\ell-1}dt)\frac{1}{r_\ell^t} + \mu_{\ell-1}dt} + o(dt).$$

Rearranging, we get

$$\frac{r_\ell^{t+dt} - r_\ell^t}{dt} = \frac{\mu_{\ell-1} - \mu_\ell - \mu_{\ell-1} r_\ell^t + \mu_\ell r_{\ell+1}^t}{(1-\mu_{\ell-1}dt)\frac{1}{r_\ell^t} + \mu_{\ell-1}dt} + o(dt)/dt.$$

Letting $dt \to 0$, we obtain

$$\dot{r}_\ell^t = r_\ell^t \left(\mu_{\ell-1} - \mu_\ell - \mu_{\ell-1} r_\ell^t + \mu_\ell r_{\ell+1}^t\right). \tag{C.19}$$

(C.19) forms a system of ordinary differential equations. The boundary condition is defined as follows. For $\ell \le \bar{K}$,

$$r_\ell^0 = \frac{\tilde{\gamma}_\ell^0}{\tilde{\gamma}_{\ell-1}^0} = \frac{p_\ell \mu_\ell}{p_{\ell-1}\mu_{\ell-1}} = \frac{\tilde{\lambda}_{\ell-1}}{\mu_{\ell-1}}. \tag{C.20}$$

where the second equality uses the fact that $\tilde{\gamma}_\ell^0 = p_\ell \mu_\ell \backslash \sum_{i=1}^{\infty} p_i \mu_i$ for each ℓ, while the third one uses (B) whereby $\frac{p_\ell}{p_{\ell-1}} = \frac{\tilde{\lambda}_{\ell-1}}{\mu_\ell}$. It is routine to see that the system of ODEs (C.19) together with the boundary condition (C.20) admits a unique solution $(r_\ell^t)_\ell$ for all $t \ge 0$.[C.42]

We first claim that $\dot{r}_\ell^0 \le 0$ for all $\ell = 2, ..., \bar{K}$. It follows from (C.19) that $\dot{r}_\ell^0 \le 0$ if and only if

$$\mu_{\ell-1} - \mu_\ell \le \mu_{\ell-1} r_\ell^0 - \mu_\ell r_{\ell+1}^0. \tag{C.21}$$

Consider any $\ell \le \bar{K}$. Substituting (C.20) into (C.21), the condition simplifies to:

$$\mu_{\ell-1} - \mu_\ell \le \tilde{\lambda}_{\ell-1} - \tilde{\lambda}_\ell,$$

which holds by Lemma C.15.

Having established that $\dot{r}_\ell^0 \le 0$ for each $\ell = 2, ..., \bar{K}$, we next prove that $\dot{r}_\ell^t \le 0$, for all $t > 0$. To this end, suppose this is not the case. Then, there exists

$$\ell \in \arg\min_{\ell'} T_{\ell'},$$

where

$$T_{\ell'} \triangleq \inf\{t' : \dot{r}_{\ell'}^{t'} > 0\}$$

if the infimum is well defined, or else $T_{\ell'} \triangleq \infty$. Let $t = T_\ell < \infty$, by the hypothesis. Then, we must have

$$\ddot{r}_\ell^t > 0; \dot{r}_{\ell'}^t \le 0, \forall \ell' \ne \ell; \text{ and } \dot{r}_\ell^t = 0.$$

Differentiating (C.19) on both sides, we obtain

$$0 < \ddot{r}_\ell^t = \dot{r}_\ell^t \left(\mu_{\ell-1} - \mu_\ell - \mu_{\ell-1} r_\ell^t + \mu_\ell r_{\ell+1}^t\right) - r_\ell^t(\mu_{\ell-1}\dot{r}_\ell^t - \mu_\ell \dot{r}_{\ell+1}^t) = r_\ell^t \mu_\ell \dot{r}_{\ell+1}^t \le 0,$$

a contradiction. We thus conclude that $\dot{r}_\ell^t \le 0$, for all $\ell = 2, ..., \bar{K}$, for all $t \ge 0$.

[C.42] This follows from the observation that the RHS of (C.19) is locally Lipschitzian in r (a fact implied by the continuous differentiability of RHS in r_ℓ^t's). See Hale p. 18, Theorem 3.1, for instance.

C.3.2 The case of $\bar{K} = \infty$.

We first derive the infinite system of ODEs in terms of agents' belief of occupying queue position $k = 1, ..., \infty$ at time t. It follows from (2), together with $q_i^* = \mu_i - \mu_{i-1}$, that

$$\tilde{\gamma}_\ell^{t+dt} = \frac{(1-\mu_\ell dt)\tilde{\gamma}_\ell^t + \mu_\ell dt \tilde{\gamma}_{\ell+1}^t}{\sum_{i=1}^{\bar{K}} \tilde{\gamma}_i^t (1 - q_i^* dt)} + o(dt).$$

$$\begin{aligned}
\frac{\tilde{\gamma}_k^{t+dt} - \tilde{\gamma}_k^t}{dt} &= \frac{(1-\mu_k dt)\tilde{\gamma}_k^t + \mu_k dt \tilde{\gamma}_{k+1}^t}{dt \sum_{i=1}^{\infty} \tilde{\gamma}_i^t (1 - q_i^* dt)} - \frac{\tilde{\gamma}_k^t}{dt} + \frac{o(dt)}{dt} \\
&= \frac{(1-\mu_k dt)\tilde{\gamma}_k^t + \mu_k dt \tilde{\gamma}_{k+1}^t}{dt \left[1 - dt \sum_{i=1}^{\infty} \tilde{\gamma}_i^t q_i^*\right]} - \frac{\tilde{\gamma}_k^t}{dt} + \frac{o(dt)}{dt} \\
&= \frac{(1-\mu_k dt)\tilde{\gamma}_k^t + \mu_k dt \tilde{\gamma}_{k+1}^t}{dt \left[1 - dt \sum_{i=1}^{\infty} \tilde{\gamma}_i^t (\mu_i - \mu_{i-1})\right]} - \frac{\tilde{\gamma}_k^t}{dt} + \frac{o(dt)}{dt} \\
&= \frac{(1-\mu_k dt)\tilde{\gamma}_k^t + \mu_k dt \tilde{\gamma}_{k+1}^t - \tilde{\gamma}_k^t \left[1 - dt \sum_{i=1}^{\infty} \tilde{\gamma}_i^t (\mu_i - \mu_{i-1})\right]}{dt \left[1 - dt \sum_{i=1}^{\infty} \tilde{\gamma}_i^t (\mu_i - \mu_{i-1})\right]} + \frac{o(dt)}{dt} \\
&= \frac{-\mu_k \tilde{\gamma}_k^t + \mu_k \tilde{\gamma}_{k+1}^t + \tilde{\gamma}_k^t \left[\sum_{i=1}^{\infty} \tilde{\gamma}_i^t (\mu_i - \mu_{i-1})\right]}{\left[1 - dt \sum_{i=1}^{\infty} \tilde{\gamma}_i^t (\mu_i - \mu_{i-1})\right]} + \frac{o(dt)}{dt}.
\end{aligned}$$

Letting $dt \to 0$, we obtain: for all $k \in \mathbb{N}$,

$$\dot{\tilde{\gamma}}_k^t = -\mu_k \tilde{\gamma}_k^t + \mu_k \tilde{\gamma}_{k+1}^t + \tilde{\gamma}_k^t \left[\sum_{i=1}^{\infty} \tilde{\gamma}_i^t (\mu_i - \mu_{i-1})\right] \triangleq f_k(\tilde{\gamma}^t), \tag{C.22}$$

and let $f \triangleq (f_k)_{k \in \mathbb{N}}$.

Let $\mathbf{X}$ be the set of sequences in ℓ^1-space endowed with ℓ^1-norm. As is well-known, this is a Banach space. Clearly, $\Delta(\mathbb{Z}_+) \subseteq \mathbf{X}$. Further, we can see that f maps from $\mathbf{X}$ to $\mathbf{X}$. Indeed, for any $\tilde{\gamma}^t \in \mathbf{X}$, we have $f(\tilde{\gamma}^t) \in \mathbf{X}$:

$$\begin{aligned}
||f(\tilde{\gamma}^t)|| &= \sum_{k=1}^{\infty} \left|f_k(\tilde{\gamma}_k^t)\right| \\
&= \sum_{k=1}^{\infty} \left| -\mu_k \tilde{\gamma}_k^t + \mu_k \tilde{\gamma}_{k+1}^t + \tilde{\gamma}_k^t \left[\sum_{i=1}^{\infty} \tilde{\gamma}_i^t (\mu_i - \mu_{i-1})\right]\right| \\
&\leq \sum_{k=1}^{\infty} \left|-\mu_k \tilde{\gamma}_k^t\right| + \sum_{k=1}^{\infty} \left|\mu_k \tilde{\gamma}_{k+1}^t\right| + \sum_{k=1}^{\infty} \left|\tilde{\gamma}_k^t \left[\sum_{i=1}^{\infty} \tilde{\gamma}_i^t (\mu_i - \mu_{i-1})\right]\right| \\
&\leq \bar{\mu} \sum_{k=1}^{\infty} \left|\tilde{\gamma}_k^t\right| + \bar{\mu} \sum_{k=1}^{\infty} \left|\tilde{\gamma}_{k+1}^t\right| + \bar{\mu} \left(\sum_{k=1}^{\infty} \left|\tilde{\gamma}_k^t\right|\right) \left(\sum_{i=1}^{\infty} \left|\tilde{\gamma}_i^t\right|\right) < \infty
\end{aligned}$$

where we recall that $\bar{\mu} \triangleq \sup_k \mu_k < \infty$ and use the fact that $\tilde{\gamma}^t \in \mathbf{X}$.

Next, consider the restriction of f defined as follows $f : U \to \mathbf{X}$ where $U \triangleq \{\{x_k\}_{k\geq 1} \in \mathbf{X} : \sum_{k=1}^{\infty} |x_k| < 1+\varepsilon\} \subset \mathbf{X}$, for some $\varepsilon > 0$, is an open set containing $\Delta(\mathbb{N})$. We show that f (restricted to U) is Lipschitz continuous. Indeed, for any $\tilde{\gamma}$ and $\tilde{\gamma}'$ in U,

$$\begin{aligned}
\|f(\tilde{\gamma}') - f(\tilde{\gamma})\| &= \sum_{k=1}^{\infty} |f_k(\tilde{\gamma}') - f_k(\tilde{\gamma})| \\
&\leq \sum_{k=1}^{\infty} |-\mu_k \tilde{\gamma}'_k + \mu_k \tilde{\gamma}_k| + \sum_{k=1}^{\infty} |\mu_k \tilde{\gamma}'_{k+1} - \mu_k \tilde{\gamma}_{k+1}| \\
&\quad + \sum_{k=1}^{\infty} \left| \tilde{\gamma}'_k \left[\sum_{i=1}^{\infty} \tilde{\gamma}'_i (\mu_i - \mu_{i-1}) \right] - \tilde{\gamma}_k \left[\sum_{i=1}^{\infty} \tilde{\gamma}_i (\mu_i - \mu_{i-1}) \right] \right| \\
&\leq \sum_{k=1}^{\infty} \mu_k |\tilde{\gamma}'_k - \tilde{\gamma}_k| + \sum_{k=1}^{\infty} \mu_k |\tilde{\gamma}'_{k+1} - \tilde{\gamma}_{k+1}| \\
&\quad + \max\{ \left[\sum_{i=1}^{\infty} |\tilde{\gamma}'_i| (\mu_i - \mu_{i-1}) \right], \left[\sum_{i=1}^{\infty} |\tilde{\gamma}_i| (\mu_i - \mu_{i-1}) \right] \} \sum_{k=1}^{\infty} |\tilde{\gamma}'_k - \tilde{\gamma}_k| \\
&\leq \bar{\mu} \|\tilde{\gamma}' - \tilde{\gamma}\| + \bar{\mu} \|\tilde{\gamma}' - \tilde{\gamma}\| + (1+\varepsilon)\bar{\mu} \|\tilde{\gamma}' - \tilde{\gamma}\|.
\end{aligned}$$

Thus, f restricted to U is Lipschitz continuous with Lipschitz constant equal to $\bar{\mu}(3+\varepsilon)$. This fact is useful for two purposes. First, it allows us to get existence and uniqueness of a solution to the system of ODEs given by (C.22) with an initial condition in $\Delta(\mathbb{Z}_+)$. In order to see this, let us consider the system of ODEs given by (C.22) where the vector field f is the mapping from $\mathbf{X}$ to $\mathbf{X}$. Since f is bounded and Lipschitz continuous on $\Delta(\mathbb{Z}_+)$ and $\Delta(\mathbb{Z}_+)$ is positively invariant, existence and uniqueness of a solution for our system of ODEs with initial condition in $\Delta(\mathbb{Z}_+)$ follows from Picard-Lindelöf Theorem on Banach spaces.[C.43]

Second, Lipschitz continuity of f restricted to U also gives us Grönwall's inequality (on continuous dependence of solutions to initial conditions) that we will use later on.

We let $\tilde{\gamma}^K(t) = (\tilde{\gamma}^K_k(t))_{k\in\mathbb{Z}_+}$ denote a solution to the system given by (C.22)

$$\dot{\tilde{\gamma}}^t = f(\tilde{\gamma}^t),$$

when $\tilde{\gamma}^0 = \tilde{\gamma}^K(0) \triangleq (\tilde{\gamma}^K_k(0))_{k\in\mathbb{Z}_+}$ where $\tilde{\gamma}^K_k(0)$ is an agent's belief of entering the queue with position k at $t = 0$ under the cutoff policy with $K < \infty$ without any randomization.[C.44] Meanwhile, $\tilde{\gamma}^\infty(t) = (\tilde{\gamma}^\infty_k(t))_{k\in\mathbb{Z}_+}$ denotes a solution to this system of ODEs when $\tilde{\gamma}^0 = \tilde{\gamma}^\infty(0) \triangleq (\tilde{\gamma}^\infty_k(0))_{k\in\mathbb{Z}_+}$ where $\tilde{\gamma}^\infty_k(0)$ is an agent's belief of entering the queue with position k at $t = 0$ under the cutoff policy with $K = \infty$.

We now complete the proof of Lemma 3 when $\bar{K} = \infty$ with the following steps.

[C.43] Recall that a subset S of $\mathbf{X}$ is positively invariant if no solution starting inside S can leave S in the future.

[C.44] Note that $\tilde{\gamma}^K_k(0) = 0$ for all $k > K$.

Step 1. $\left\|\tilde{\gamma}^K(0) - \tilde{\gamma}^\infty(0)\right\| \to 0$ as $K \to \infty$.

Proof. We know that for all $\ell = 2, ..., K$: $\tilde{\gamma}_\ell^K(0) = \prod_{i=2}^{\ell} r_i^0 \tilde{\gamma}_1^K(0) = \prod_{i=2}^{\ell} \frac{\lambda_{i-1}}{\mu_{i-1}} \tilde{\gamma}_1^K(0)$, where we used (C.20), while $\tilde{\gamma}_\ell^K(0) = 0$ for $\ell \geq K+1$. In addition, we know that $\tilde{\gamma}_\ell^\infty(0) = \prod_{i=2}^{\ell} \frac{\tilde{\lambda}_{i-1}}{\mu_{i-1}} \tilde{\gamma}_1^\infty(0)$ and $\sum_{k=1}^{\infty} \prod_{i=2}^{k} \frac{\tilde{\lambda}_{i-1}}{\mu_{i-1}} \tilde{\gamma}_1^\infty(0) = 1$ where our convention is that $\prod_{i=2}^{1} \frac{\tilde{\lambda}_{i-1}}{\mu_{i-1}} \triangleq 1$. Thus,

$$\tilde{\gamma}_1^\infty(0) = \frac{1}{\sum_{k=1}^{\infty} \prod_{i=2}^{k} \frac{\tilde{\lambda}_{i-1}}{\mu_{i-1}}}.$$

Note that this implies that $\sum_{k=1}^{\infty} \prod_{i=2}^{k} \frac{\tilde{\lambda}_{i-1}}{\mu_{i-1}} < \infty$. Similar computation yields

$$\tilde{\gamma}_1^K(0) = \frac{1}{\sum_{k=1}^{K} \prod_{i=2}^{k} \frac{\tilde{\lambda}_{i-1}}{\mu_{i-1}}}.$$

Note that $\left|\tilde{\gamma}_1^K(0) - \tilde{\gamma}_1^\infty(0)\right| \to 0$ as K increases. We have

$$\begin{aligned}
\left\|\tilde{\gamma}^K(0) - \tilde{\gamma}^\infty(0)\right\| &= \sum_{k=1}^{\infty} \left|\tilde{\gamma}_k^K(0) - \tilde{\gamma}_k^\infty(0)\right| \\
&= \sum_{k=1}^{K} \left|\tilde{\gamma}_k^K(0) - \tilde{\gamma}_k^\infty(0)\right| + \sum_{k=K+1}^{\infty} \left|\tilde{\gamma}_k^\infty(0)\right| \\
&= \sum_{k=1}^{K} \left| \prod_{i=2}^{k} \frac{\tilde{\lambda}_{i-1}}{\mu_{i-1}} \tilde{\gamma}_1^K(0) - \prod_{i=2}^{k} \frac{\tilde{\lambda}_{i-1}}{\mu_{i-1}} \tilde{\gamma}_1^\infty(0) \right| + \sum_{k=K+1}^{\infty} \left|\tilde{\gamma}_k^\infty(0)\right| \\
&= \left|\tilde{\gamma}_1^K(0) - \tilde{\gamma}_1^\infty(0)\right| \sum_{k=1}^{K} \prod_{i=2}^{k} \frac{\tilde{\lambda}_{i-1}}{\mu_{i-1}} + \sum_{k=K+1}^{\infty} \left|\tilde{\gamma}_k^\infty(0)\right|.
\end{aligned}$$

Since $\left|\tilde{\gamma}_1^K(0) - \tilde{\gamma}_1^\infty(0)\right| \to 0$ as $K \to \infty$, $\sum_{k=1}^{\infty} \prod_{i=2}^{k} \frac{\tilde{\lambda}_{i-1}}{\mu_{i-1}} < \infty$, and $\sum_{k=K+1}^{\infty} \left|\tilde{\gamma}_k^\infty(0)\right|$ goes to 0 as $K \to \infty$, the result follows. ∎

Step 2. For each $t > 0$,

$$\lim_{K \to \infty} \sum_{k=1}^{\infty} \left|\tilde{\gamma}_k^K(t) - \tilde{\gamma}_k^\infty(t)\right| = 0.$$

Proof. By Grönwall's inequality,

$$\left\|\tilde{\gamma}^K(t) - \tilde{\gamma}^\infty(t)\right\| \leq e^{Ct} \left\|\tilde{\gamma}^K(0) - \tilde{\gamma}^\infty(0)\right\|,$$

where, as shown earlier, $C \triangleq \bar{\mu}(3 + \varepsilon)$ is the Lipschitz constant for the Lipschitz continuous function f restricted to open set $U = \{\{x_k\}_{k \geq 1} \in \mathbf{X} : \sum_{k=1}^{\infty} |x_k| < 1 + \varepsilon\}$. The result then follows from Step 1. ∎

Step 3. $\dot{r}_\ell^\infty(t) \leq 0$ for all $\ell \geq 2$ and t, where $r_\ell^\infty(t) = \tilde{\gamma}_\ell^\infty(t)/\tilde{\gamma}_{\ell-1}^\infty(t)$ for all $\ell \geq 2$.

Proof. Recall from (C.19) that the system of ODEs is given by

$$\dot{r}_\ell^\infty(t) = r_\ell^\infty(t)\left(\mu_{\ell-1} - \mu_\ell - \mu_{\ell-1} r_\ell^\infty(t) + \mu_\ell r_{\ell+1}^\infty(t)\right)$$

for all $\ell \geq 2$. Suppose to the contrary that $\dot{r}_\ell^\infty(t) > 0$ for some ℓ and t. We already proved in the previous section where $\bar{K} < \infty$ that

$$\dot{r}_\ell^K(t) = r_\ell^K(t)\left(\mu_{\ell-1} - \mu_\ell - \mu_{\ell-1} r_\ell^K(t) + \mu_\ell r_{\ell+1}^K(t)\right) \leq 0$$

for all $K < \infty$, ℓ and t. To show a contradiction, it is enough to prove that

$$r_\ell^K(t)\left(\mu_{\ell-1} - \mu_\ell - \mu_{\ell-1} r_\ell^K(t) + \mu_\ell r_{\ell+1}^K(t)\right) \to r_\ell^\infty(t)\left(\mu_{\ell-1} - \mu_\ell - \mu_{\ell-1} r_\ell^\infty(t) + \mu_\ell r_{\ell+1}^\infty(t)\right)$$

as $K \to \infty$. To this end, it suffices to show that $r_\ell^K(t)$ and $r_{\ell+1}^K(t)$ converge respectively to $r_\ell^\infty(t)$ and $r_{\ell+1}^\infty(t)$. It follows from Step 2 that for each k :

$$\lim_{K\to\infty} \tilde{\gamma}_k^K(t) = \tilde{\gamma}_k^\infty(t).$$

By assumption $\tilde{\gamma}_k^\infty(0) > 0$ for all k, so

$$\lim_{K\to\infty} r_\ell^K(t) = \lim_{K\to\infty} \frac{\tilde{\gamma}_\ell^K(t)}{\tilde{\gamma}_{\ell-1}^K(t)} = \frac{\tilde{\gamma}_\ell^\infty(t)}{\tilde{\gamma}_{\ell-1}^\infty(t)} = r_\ell^\infty(t)$$

Similarly,

$$\lim_{K\to\infty} r_{\ell+1}^K(t) = r_{\ell+1}^\infty(t),$$

which completes the argument. ∎

Step 3 produces the desired result. ∎

D Proofs from Section 6

D.1 Proofs of Proposition 2

Consider FCFS with full information. We need to show that (IC_t) holds for all $t \geq 0$. With the full information rule, we only need to show that (IC_0) holds. By Lemma 1, condition (IC_0) can be written as:

$$V - C\frac{k}{\mu_k} \geq 0 \iff \mu_k V - Ck \geq 0 \tag{D.23}$$

for all $k \leq K^*$. In the sequel, we let K^{FI} be the largest integer satisfying (D.23). We know that, by regularity of μ, $k \mapsto \mu_k V - Ck$ is single-peaked (by Lemma A.5 for $\alpha = 1$ and

$\xi = 0$). Hence, K^{FI} is well-defined (i.e., finite) given our assumption that μ_k is uniformly bounded. In addition, Equation (D.23) holds at state k if and only if $k \leq K^{FI}$. Hence, it is enough for our purpose to show that $K^* \leq K^{FI}$.

Proceed by contradiction and assume that the optimal cutoff policy p^*, which we recall solves $[P']$, puts strictly positive weight on $k > K^{FI}$. Note that, using again the fact that $k \mapsto \mu_k V - Ck$ is single-peaked, for any such k, $\mu_k V - Ck < 0$. Now, build p' such that $p'_k = 0$ for all $k > K^{FI}$ and $p'_k = Zp^*_k$ for all $k \leq K^{FI}$ where $Z > 1$ is set so that the sum of p'_k is equal to 1. Given that p^* satisfies (B$'$) and given that, by construction, $p'_k/p'_{k-1} = p^*_k/p^*_{k-1}$ for all $k \leq K^{FI}$, we must have that p' also satisfies (B$'$). Compared to p^*, distribution p' removes all weight on negative values and, for each positive value, increases its weight. This must strictly increase the value of the objective. It remains to show that p' satisfies (IR). The value of the objective must be positive under p^* (recall that the dirac mass on 0 brings a value of the objective of 0), and so the value of the objective must be positive under p' as well. Given that $\alpha = 1$, this implies that (IR) is satisfied.

Now, we show that the same result holds for SIRO with full information. We fix the optimal cutoff outcome (x^*, y^*, p^*) and recall that K^* is finite since $K^* \leq K^{FI} < \infty$. We need to show that under SIRO with full information (IC_t) holds for all $t \geq 0$. Again, with full information, we only need to show that (IC_0) holds. We denote τ^S_k (τ^F_k) the expected waiting time of an agent who joins a queue with $k-1$ agents when the queue discipline is SIRO (FCFS). Next, we claim $\tau^S_{K^*} \leq \tau^F_{K^*}$. This follows from Lemma 1 and Lemma 4. Then,

$$V - C\tau^S_{K^*} \geq V - C\tau^F_{K^*} \geq 0.$$

Given that τ^S_k is nondecreasing in k (see Lemma 4 for a formal argument), we obtain that $V - C\tau^S_k \geq 0$ for all $k \leq K^*$ which yields (IC_0) for SIRO with full information.

For the second statement, fix the outcome (x^*, y^*, p^*) associated with the optimal cutoff policy. By Little's law, the expected waiting time for an agent when he joins the queue is given by

$$T^* \triangleq \frac{\sum_k p^*_k k}{\sum_k p^*_k \mu_k},$$

where the numerator is the average length of queue and the denominator is the average exit rate (recall that $y^*_k = 0$ for all k), both evaluated at the invariant distribution p^*. Hence, if an agent is given no information except for the recommendation to join the queue, he will join the queue as long as

$$V - cT^* \geq 0 \Leftrightarrow V \sum_k p^*_k \mu_k - C \sum_k p^*_k k \geq 0,$$

which holds since p^* satisfies (IR). ∎

D.2 Proof of Proposition 3

Fix any queueing rule $q \in Q$, and suppose the information rule $I \in \mathcal{I}$ admits distinct beliefs for expected waiting times. For each belief $\gamma \in \text{supp}(I_0)$, let $\tau(\gamma)$ denote the expected waiting time. Let $\overline{\tau} \triangleq \sup_{\gamma \in \text{supp}(I)} \tau(\gamma)$. Then, again by Little's law,

$$T^* = \int_{\gamma \in \text{supp}(I)} \tau(\gamma) I(d\gamma).$$

We must have $\overline{\tau} > T^*$, or else $\tau(\gamma) = T^*$ for all $\gamma \in \text{supp}(I)$, which contradicts the hypothesis that distinct expected waiting time are possible for an agent at time $t = 0$. Since $\overline{\tau} > T^*$, there exists $\gamma \in \text{supp}(I)$ such that $\tau(\gamma) > T^*$. Then, an agent with belief γ would receive the payoff

$$V - \tau(\gamma)C < V - T^*C = 0$$

by obeying the designer's recommendation to join the queue at $t = 0$. Hence, (IC_0) fails under (q, I). ∎

D.3 Proof of Lemma 4

The agent in state $k \geq 1$ waits for

$$\tau_k = (\frac{\mu_k}{k} dt)dt + \lambda_k dt(dt + \tau_{k+1}) + [1 - \mu_k dt - \lambda_k dt] (\tau_k + dt) + \frac{k-1}{k} \mu_k dt(\tau_{k-1} + dt) + o(dt),$$

where $\lambda_{K^*} = 0$.

Simplifying and letting $dt \to 0$, we get

$$\tau_k = \frac{1}{\mu_k + \lambda_k} + \frac{\lambda_k}{\mu_k + \lambda_k} \tau_{k+1} + \frac{k-1}{k} \frac{\mu_k}{\mu_k + \lambda_k} \tau_{k-1}. \tag{D.24}$$

Let $\Phi(\tau)$ denote the vector-valued map defined by the RHS of the above equations. Note that

$$||\Phi(\tau'') - \Phi(\tau')|| \leq \xi ||\tau'' - \tau'||,$$

where $\xi \triangleq \max_k \{\frac{\lambda_k}{\mu_k + \lambda_k} + \frac{k-1}{k} \frac{\mu_k}{\mu_k + \lambda_k}\} < 1$.[D.45] Hence, the self-map $\Phi : \mathbb{R}_+^{K^*} \to \mathbb{R}_+^{K^*}$ is a contraction, so admits a unique fixed point. Therefore, we conclude that τ is well defined.

We next prove that τ_k is nondecreasing in k; the proof consists of several steps.

Step 1. If $\tau_k < \frac{k}{\mu_k}$, then $k > 1$ and $\tau_k > \tau_{k-1}$.

Proof. Suppose to the contrary that $\tau_k < \frac{k}{\mu_k}$, and either $k = 1$ or $\tau_k \leq \tau_{k-1}$. Then, use (D.24) to write

$$\lambda_k(\tau_{k+1} - \tau_k) = \frac{\mu_k}{k} \tau_k - 1 + \frac{k-1}{k} \mu_k (\tau_k - \tau_{k-1}) < 0, \tag{D.25}$$

[D.45] The norm $||\cdot||$ denotes the norm max.

since $\tau_k < \frac{k}{\mu_k}$ and either $k = 1$ or $\tau_k \le \tau_{k-1}$. It then follows that

$$\tau_{k+1} < \tau_k < \frac{k}{\mu_k} \le \frac{k+1}{\mu_{k+1}},$$

so one can iterate the argument repeatedly. Eventually, we obtain $\tau_{K^*} < \frac{K^*}{\mu_{K^*}}$ and $\tau_{K^*} < \tau_{K^*-1}$. But (D.25) with $k = K^*$ and $\lambda_k = 0$ yields

$$\frac{K^*-1}{K^*}\mu_{K^*}(\tau_{K^*} - \tau_{K^*-1}) = 1 - \frac{\mu_{K^*}}{K^*}\tau_{K^*} > 0, \tag{D.26}$$

since $\tau_{K^*} < \frac{K^*}{\mu_{K^*}}$. Since this implies that $\tau_{K^*-1} < \tau_{K^*}$, we have a contradiction. ∎

Since $\tau_1 \ge \frac{1}{\mu_1}$ by Step 1 and, by assumption, $K^* < \infty$, the states $\{1, ..., K^*\}$ are partitioned into a finite set of consecutive intervals $\mathcal{K}_1, ..., \mathcal{K}_m$, where $\mathcal{K}_i = \{k_i, k_i + 1, ..., k_i + j_i\}$, $\mathcal{K}_{i+1} = \{k_i + j_i + 1, ..., k_i + j_i + j_{i+1}\}$, for some $k_i, j_i, j_{i+1} > 0$, such that, $\tau_k \ge \frac{k}{\mu_k}$ for each $k \in \mathcal{K}_i$, if i is an odd number, and $\tau_k < \frac{k}{\mu_k}$ for each $k \in \mathcal{K}_i$, if i is an even number.

Step 2. Suppose $k \in \mathcal{K}_i$, where i is an odd number. Then, $\tau_{k+1} \ge \tau_k$. In addition, $\tau_{k_i} > \tau_{k_i-1}$ when $i > 1$.

Proof. Let $\mathcal{K}_i = \{k_i, k_i + 1, ..., k_i + j_i\}$. First, let $k = k_i$. There are two cases. Suppose first $i = 1$, so $k_i = 1$. In this case, (D.25) with $k = 1$ gives

$$\lambda_1(\tau_2 - \tau_1) = \mu_1\tau_1 - 1 \ge 0,$$

where the inequality follows from Step 1. So $\tau_2 \ge \tau_1$. Suppose next $i > 1$. Then, we have

$$\tau_{k_i} \ge \frac{k_i}{\mu_{k_i}} \ge \frac{k_i - 1}{\mu_{k_i-1}} > \tau_{k_i-1},$$

where the last inequality follows from the fact that $k_i - 1 \in \mathcal{K}_{i-1}$, where $i - 1$ is an even number. Using this and (D.25), we conclude that for $k = k_i$,

$$\lambda_k(\tau_{k+1} - \tau_k) = \frac{\mu_k}{k}\tau_k - 1 + \frac{k-1}{k}\mu_k(\tau_k - \tau_{k-1}) \ge 0,$$

where the inequality follows since $k \in \mathcal{K}_i$ with i being odd. Hence, $\tau_{k_i+1} \ge \tau_{k_i}$. One can then iterate the argument to conclude that $\tau_{k+1} \ge \tau_k$ for any $k \in \mathcal{K}_i$. ∎

Since, by Step 1, $\tau_k > \tau_{k-1}$ for any $k \in \mathcal{K}_i$, where i is even, Step 2 implies that $\tau_{[\cdot]}$ is nonincreasing. Step 1 further yields $\tau_1 \ge 1/\mu_1$. Meanwhile, it follows from (D.26) that $\tau_{K^*} \le K^*/\mu_{K^*}$, since $\tau_{K^*} \ge \tau_{K^*-1}$.

Next, assume $1/\mu_1 < K^*/\mu_{K^*}$. We prove that $\tau_{K^*} > \tau_1$. Suppose not. Then, τ_k must be constant in k for all $k = 1, ..., K^*$. The recursion equations (D.25) then implies that $\tau_k = k/\mu_k$ for all $k = 1, ..., K^*$. But this implies that $1/\mu_1 = K^*/\mu_{K^*}$, a contradiction.

Finally, suppose $2\mu_1 > \mu_2$. We first prove $\tau_2 > \tau_1$. Suppose to the contrary that $\tau_2 = \tau_1$. Then, (D.25) with $k = 1$ implies that $\tau_1 = 1/\mu_1$. Since $\tau_3 \geq \tau_2$, (D.25) with $k = 2$ gives $\tau_2 = 2/\mu_1$. But since $2\mu_1 > \mu_2$, we have a contradiction to $\tau_2 = \tau_1$. Now that $\tau_2 > \tau_1$, the proof of Step 2 holds with strict inequality—i.e., $\tau_{k+1} > \tau_k$ for all $k \in \mathcal{K}_i$ with i being odd numbered. Combining Step 1 and Step 2 (with strict inequality) then produces strict monotonicity of τ_k in k. ∎

D.4 Proof of Proposition 4

Let γ^t be the (degenerate) posterior belief induced by this information policy. Since the expected waiting at each t depends only on the queue length at each t (and not on one's queue position), we can effectively focus on belief $\gamma^t = (\gamma^t_k)$, where γ^t_k is the posterior probability at time t that the queue length is k.

As before, we focus on the likelihood ratio $r^t_k \triangleq \frac{\gamma^t_k}{\gamma^t_{k-1}}$. For $k \geq 2$ and for a small period dt, the likelihood ratio is updated from time t to $t + dt$ according to the following recursion:

$$
\begin{aligned}
r_k^{t+dt} =& \frac{\gamma_k^t(1-\tilde{\lambda}_k dt - \mu_k dt) + \gamma_{k-1}^t \tilde{\lambda}_{k-1} dt + \gamma_{k+1}^t \mu_{k+1} \frac{k}{k+1} dt}{\gamma_{k-1}^t(1-\tilde{\lambda}_{k-1} dt - \mu_{k-1} dt) + \gamma_{k-2}^t \tilde{\lambda}_{k-2} dt + \gamma_k^t \mu_k \frac{k-1}{k} dt} + o(dt) \\
=& \frac{r_k^t(1-\tilde{\lambda}_k dt - \mu_k dt) + \tilde{\lambda}_{k-1} dt + r_{k+1}^t r_k^t \mu_{k+1} \frac{k}{k+1} dt}{(1-\tilde{\lambda}_{k-1} dt - \mu_{k-1} dt) + \frac{1}{r_{k-1}^t} \tilde{\lambda}_{k-2} dt + r_k^t \mu_k \frac{k-1}{k} dt} + o(dt),
\end{aligned}
$$

where we use the convention that $r^t_{K^*+1} = 0$. (Recall also that $\tilde{\lambda}_k$ is an effective arrival rate: $\tilde{\lambda}_k \triangleq \lambda_k$ for $k < K^* - 1$ and $\tilde{\lambda}_{K^*-1} \triangleq x^*_{K^*-1}\lambda_{K^*-1}$.) To understand the formula, consider the numerator, which corresponds to the belief that the queue length is k at time $t + dt$: it equals γ^t_k, the corresponding belief at time t, if either no agent arrives or no agent is served, γ^t_{k-1} if an agent arrives during $[t, t + dt)$, and γ^t_{k+1} if an agent other than himself is served during $[t, t + dt)$ whose probability is $\frac{k}{k+1}\mu_{k+1}dt$. The denominator, which corresponds to the belief about the queue length being $k - 1$ at time $t + dt$, is explained analogously.

Subtracting r^t_k from both sides and dividing it by dt and let $dt \to 0$, we obtain:

$$
\dot{r}_k^t = r_k^t \left[\mu_{k-1} + \tilde{\lambda}_{k-1} - \mu_k - \tilde{\lambda}_k + \frac{\tilde{\lambda}_{k-1}}{r_k^t} + \mu_{k+1} r_{k+1}^t \frac{k}{k+1} - \frac{\tilde{\lambda}_{k-2}}{r_{k-1}^t} - \mu_k r_k^t \frac{k-1}{k} \right].
$$

We now evaluate this at $t = 0$. As before we use $r^0_k = \frac{\tilde{\lambda}_{k-1}}{\mu_{k-1}}$ to obtain: for each $k = 2, ..., K^*$,

$$
\begin{aligned}
\dot{r}_k^0 &= r_k^0 \left[\mu_{k-1} + \tilde{\lambda}_{k-1} - \mu_k - \tilde{\lambda}_k + \mu_{k-1} + \tilde{\lambda}_k \frac{\mu_{k+1}}{\mu_k} \frac{k}{k+1} - \mu_{k-2} - \tilde{\lambda}_{k-1} \frac{\mu_k}{\mu_{k-1}} \frac{k-1}{k} \right] \\
&= r_k^0 \left[2\mu_{k-1} - \mu_k - \mu_{k-2} + \tilde{\lambda}_{k-1} \left(1 - \frac{\mu_k}{\mu_{k-1}} \frac{k-1}{k} \right) - \tilde{\lambda}_k \left(1 - \frac{\mu_{k+1}}{\mu_k} \frac{k}{k+1} \right) \right], \qquad \text{(D.27)}
\end{aligned}
$$

where $\mu_0 \equiv 0$ and $\tilde{\lambda}_{K^*} \equiv 0$.

We first claim that $\dot{r}_k^0 > 0$ for all $k = 2, ..., K^*$. This follows for $k = 2, ..., K^* - 2$ since the RHS of (D.27) has the same sign as:

$$\begin{aligned}
&2\mu_{k-1} - \mu_k - \mu_{k-2} + \lambda_{k-1}\left(1 - \frac{\mu_k}{\mu_{k-1}}\frac{k-1}{k}\right) - \lambda_k\left(1 - \frac{\mu_{k+1}}{\mu_k}\frac{k}{k+1}\right) \\
\geq &\lambda_{k-1}\left(1 - \frac{\mu_k}{\mu_{k-1}}\frac{k-1}{k}\right) - \lambda_k\left(1 - \frac{\mu_{k+1}}{\mu_k}\frac{k}{k+1}\right) > 0,
\end{aligned}$$

where the first inequality follows from the regularity of service process and the last inequality follows from the hypothesis assumed. For $k = K^* - 1$, the relevant expression is:

$$\begin{aligned}
&2\mu_{K^*-2} - \mu_{K^*-1} - \mu_{K^*-3} + \lambda_{K^*-2}\left(1 - \frac{\mu_{K^*-1}}{\mu_{K^*-2}}\frac{K^*-2}{K^*-1}\right) - x^*_{K^*-1}\lambda_{K^*-1}\left(1 - \frac{\mu_{K^*}}{\mu_{K^*-1}}\frac{K^*-1}{K^*}\right) \\
\geq &\lambda_{K^*-2}\left(1 - \frac{\mu_{K^*-1}}{\mu_{K^*-2}}\frac{K^*-2}{K^*-1}\right) - \lambda_{K^*-1}\left(1 - \frac{\mu_{K^*}}{\mu_{K^*-1}}\frac{K^*-1}{K^*}\right) > 0,
\end{aligned}$$

where again the first inequality follows from the regularity of service process (which also implies $\frac{K^*}{\mu_{K^*}} \geq \frac{K^*-1}{\mu_{K^*-1}}$; see Lemma 1). Finally, for $k = K^*$, the expression reduces to

$$2\mu_{K^*-1} - \mu_{K^*} - \mu_{K^*-2} + x^*_{K^*-1}\lambda_{K^*-1}\left(1 - \frac{\mu_{K^*}}{\mu_{K^*-1}}\frac{K^*-1}{K^*}\right) > 0,$$

since $x^*_{K^*-1} > 0$.

Since $\dot{r}_k^0 > 0$ for all $k = 2, ..., K^*$, for $t > 0$ small enough, we have $r_k^t > r_k^0$, for all $k = 2, ..., K^*$. This means that γ^t strictly dominates γ^t in FOSD. Hence, $\sum_{k=1}^{K^*} \gamma_k^t \tau_k > \sum_{k=1}^{K^*} \gamma_k^0 \tau_k$, since $\tau_{K^*} > \tau_1$ from Lemma 4. Since (IR) binds at $t = 0$, we have

$$0 = V - C\sum_{k=1}^{K^*} \gamma_k^0 \tau_k > V - C\sum_{k=1}^{K^*} \gamma_k^t \tau_k,$$

which implies that (IC_t) is violated at $t > 0$. ∎

E Proof of Proposition 5

Fix any π and no information. Beliefs evolve during $[t, t + dt)$ according to Bayes rule as follows:

$$\gamma_{1,1}^{t+dt} = \frac{\gamma_{1,1}^t(1 - \lambda dt - \mu dt) + \gamma_{2,2}^t \mu dt}{\gamma_{1,1}^t(1 - \mu dt) + \gamma_{2,1}^t(1 - \mu dt) + \gamma_{2,2}^t} + o(dt),$$

$$\gamma_{2,1}^{t+dt} = \frac{\gamma_{2,1}^t(1 - \mu dt) + \gamma_{1,1}^t \pi\lambda dt}{\gamma_{1,1}^t(1 - \mu dt) + \gamma_{2,1}^t(1 - \mu dt) + \gamma_{2,2}^t} + o(dt),$$

$$\gamma_{2,2}^{t+dt} = \frac{\gamma_{1,1}^t \lambda dt(1-\pi) + \gamma_{2,2}^t(1-\mu dt)}{\gamma_{1,1}^t(1-\mu dt) + \gamma_{2,1}^t(1-\mu dt) + \gamma_{2,2}^t} + o(dt).$$

From these, we can derive ODEs that describe belief evolutions:

$$\dot{\gamma}_{1,1}^t = -\gamma_{1,1}^t \lambda + \gamma_{2,2}^t(1-\gamma_{1,1}^t)\mu, \tag{E.28}$$

$$\dot{\gamma}_{2,1}^t = \lambda \gamma_{1,1}^t \pi - \mu \gamma_{2,1}^t \gamma_{2,2}^t, \tag{E.29}$$

$$\dot{\gamma}_{2,2}^t = \gamma_{1,1}^t \lambda(1-\pi) - (\gamma_{2,2}^t)^2 \mu, \tag{E.30}$$

with a boundary condition at $t = 0$:

$$\gamma_{1,1}^0 = \frac{p_0 \lambda}{\lambda p_0 + \lambda p_1} = \frac{\mu}{\mu+\lambda}, \gamma_{2,1}^0 = \frac{p_1 \lambda(1-\pi)}{\lambda p_0 + \lambda p_1} = \frac{\lambda(1-\pi)}{\mu+\lambda}, \gamma_{2,2}^0 = \frac{p_1 \lambda \pi}{\lambda p_0 + \lambda p_1} = \frac{\lambda \pi}{\mu+\lambda},$$

where we used the fact that $p_1 = \frac{\lambda}{\mu} p_0$ and $p_2 = (\frac{\lambda}{\mu})^2 p_0$ at the invariant distribution of k.

Substituting, we have at $t = 0$.

$$\dot{\gamma}_{1,1}^0 = -\gamma_{1,1}^0 \lambda + \gamma_{2,2}^0(1-\gamma_{1,1}^0)\mu = -\frac{\mu\lambda(\mu+(1-\pi)\lambda)}{(\mu+\lambda)^2},$$

$$\dot{\gamma}_{2,1}^0 = \lambda \gamma_{1,1}^0 \pi - \mu \gamma_{2,1}^0 \gamma_{2,2}^0 = \frac{\mu\lambda\pi(\mu+\pi\lambda)}{(\mu+\lambda)^2}.$$

Then, the change in expected waiting time at $t = 0$ is:

$$\begin{aligned}
\omega'(0) =& \dot{\gamma}_{1,1}^0 \tau_{1,1} + \dot{\gamma}_{2,1}^0 \tau_{2,1} + \dot{\gamma}_{2,2}^0 \tau_{2,2} \\
=& \dot{\gamma}_{1,1}^0 \tau_{1,1} + \dot{\gamma}_{2,1}^0 \frac{1}{\mu} + \dot{\gamma}_{2,2}^0 (\tau_{1,1} + \frac{1}{\mu}) \\
=& -\dot{\gamma}_{2,1}^0 \tau_{1,1} - \dot{\gamma}_{1,1}^0 \frac{1}{\mu} \\
=& -\frac{\mu\lambda\pi(\mu+\pi\lambda)}{(\mu+\lambda)^2} \frac{\mu+\lambda}{\mu(\mu+\lambda\pi)} + \frac{\lambda(\mu+(1-\pi)\lambda)}{(\mu+\lambda)^2} \\
=& -\frac{\lambda\pi(\mu+\lambda)}{(\mu+\lambda)^2} + \frac{\lambda(\mu+(1-\pi)\lambda)}{(\mu+\lambda)^2} \\
=& -\frac{\lambda}{(\mu+\lambda)^2} \left[\pi(\mu+\lambda) - (\mu+(1-\pi)\lambda)\right],
\end{aligned}$$

which is strictly positive whenever $\pi < \frac{\mu+\lambda}{\mu+2\lambda}$. Then, since (IR) is binding at the optimal policy, this shows that (IC_t) fails for sufficiently small $t > 0$ for any such π. ∎