Optimal Congestion Signaling under Customer Heterogeneity and Demand Variation

Jing Dong
Decision, Risk, and Operations, Columbia Business School, jing.dong@gsb.columbia.edu

Prakirt Raj Jhunjhunwala
Decision, Risk, and Operations, Columbia Business School, prj2122@columbia.edu

Yash Kanoria
Decision, Risk, and Operations, Columbia Business School, ykanoria@gmail.com

Abstract. In an unobservable queue, where customers cannot observe the queue length, a server aiming to maximize throughput benefits by providing coarse congestion information to influence customers' decisions on whether to join the queue. We characterize the structure of the server's optimal congestion signaling mechanism by adopting a Bayesian persuasion framework to model the customers' response to the provided signal. We model two important real-world features, introducing additional information asymmetries between servers and customers. Firstly, we explore customer heterogeneity with private customers' willingness to wait. Secondly, we consider demand variation, where the arrival rate of customers varies from one period to the next and is only observable to the server. We present a novel Constrained Markov Decision Process (C-MDP) formulation of the server's throughput maximization problem under the resulting customer equilibrium.

Customer heterogeneity with private customer types works in favor of the customers, allowing some customer classes to receive strictly positive utility. We show that the optimal signaling mechanism has a *laminar* structure. Counter-intuitive phenomena arise within this structure, for example, with two customer types, the interval of queue lengths in which only the more patient customers enter can be nested within the intervals of queue length in which all customers enter. In contrast to customer heterogeneity, demand variation favors the server by providing additional flexibility in signaling due to the customer not knowing the realized arrival rate. Using simulation, we compare different signaling mechanisms (including the *no* information, the *full* information, and the optimal mechanisms) and demonstrate their impact on the overall system performance.

Key words: Signaling mechanism, Delay announcement, Constrained-MDP, Laminar structure

1. Introduction

In many service systems, when a customer arrives, they receive a wait time indication from the server, which can significantly influence their decision to enter the system. For example, in hospital emergency departments, wait time indications are common and have a notable impact on patient behavior (see, e.g., Walker et al. 2022). In call centers, customers are informed about the congestion with messages like "Your call will be connected soon." or "Please hold. Our agents are currently busy." Similarly, ride-hailing services like Uber or Lyft give users wait time estimates before they book a ride. In these contexts, the server possesses accurate information about wait times, while customers lack this knowledge and rely on the service provider for this information to make informed decisions about whether to join the queue. This information asymmetry leads to a strategic interaction between the server and the customer, who have misaligned objectives; the server may seek to maximize throughput (or revenue), while the customers aim to maximize their utility by balancing the benefits of receiving service against the costs of waiting. The server, in pursuit of maximizing its own objective, can influence customer decisions to enter the queue by strategically designing wait time information at varying levels of detail. At one extreme, the server can choose to disclose *full information*, such as the current queue length. At the other extreme, it may opt to provide *no information* at all. Alternatively, the server can provide a *coarse* congestion signal, such as labeling conditions as "high congestion" or "low congestion". While these signals do not offer precise wait times if appropriately designed, they can serve as a tool to nudge customer behavior: encouraging entry with a low congestion signal and discouraging it with a high congestion signal.

Moreover, in real-world systems, there are two additional sources of information asymmetry between the server and customers that can significantly impact the server's congestion signaling decision. The first is unpredictable fluctuations in customer arrival rates (e.g., from one day to the next), with the realized arrival rate visible to the server but not to the customers. We call this *demand variation*. The second is *customer heterogeneity* in willingness to wait, which is private to each individual customer.

In this work, we ask the following question: *What congestion signaling mechanism should a throughput-maximizing server adopt? How do demand variation and customer heterogeneity in willingness to wait influence the optimal signaling strategy?*

Throughput is often the primary metric of interest since revenue is proportional to throughput in settings with a fixed price. Maximizing the server's throughput is equivalent to minimizing server idleness or maximizing the utilization rate of the server. We apply a *Bayesian persuasion* framework to study the optimal signaling mechanism for maximizing server throughput under varying demand and heterogeneous customers with private types.

Bayesian persuasion framework: The Bayesian persuasion framework (Kamenica and Gentzkow 2011) assumes that the sender (server) *commits* to a signaling mechanism revealed to all customers a priori. This allows the receivers of signals (customers), who are assumed to act as Bayesian decision-makers, to respond to a given congestion signaling mechanism and the anticipated strategies of other customers. They do this by computing the Bayes posterior of the expected wait time and then choosing their best response using these posterior beliefs. Customers are further assumed to reason based on the steady-state distribution of queue length under each arrival rate, which is endogenously determined by both the server's signaling mechanism and the customers' strategies. We study the resulting customer equilibrium, which depends on the server's chosen signaling mechanism.

In the queueing setting, customers' actions generate externalities: a customer's decision to enter or not enter the queue impacts the congestion level in the system, thereby influencing the experiences and decisions of subsequent customers. This interdependence poses a significant challenge compared to the classic

Bayesian persuasion framework, where customers' prior beliefs are exogenous and not influenced by the actions of others.

The assumption in our theoretical model that customers can estimate the relevant posterior expected queue lengths is motivated by the expectation that customers observe the empirical distribution of queue lengths they encounter under each possible signal and adapt their entry strategies accordingly. Importantly, the conditional distribution of queue lengths given a particular signal serves as a sufficient statistic for the customers' expected wait times.

Next, we present our consideration of the two features: customer heterogeneity with private types (or private patience information) and demand variation.

Customer heterogeneity with private type: Customers often vary in their patience or willingness to wait. We study a setting with K customer types, each possessing varying levels of patience, where a customer's patience is their private information. The service provider only knows the distribution of customers' patience. Such a model captures real-world applications where congestion signals are public or cannot be personalized to the customer (e.g., for legal or fairness reasons).

For a throughput-maximizing server, accounting for heterogeneity in customer patience is crucial when designing the signaling mechanism. With strategic Bayesian customers, each congestion signal encourages only customers with patience exceeding a certain cutoff level to enter while dissuading less patient customers from entering. This cutoff level, communicated through the server's signal, depends on the current congestion level (system state) and the signaling policy. As we will demonstrate, the server must carefully design the signaling mechanism to achieve a throughput-maximizing customer equilibrium: ensuring that the right set of customers, based on their patience, chooses to enter the queue at each queue length.

Demand variation: In many real-world service systems, customers' arrival rate (i.e., demand) can fluctuate significantly from day to day. For instance, in call centers and hospital emergency departments, it has been shown that a doubly stochastic Poisson process tends to more accurately represent the arrival process (see, e.g., Aldor-Noiman et al. 2009, Avramidis et al. 2004, Bassamboo et al. 2010, Ibrahim et al. 2016). Further, these variations are typically known to or relatively easy to infer by the server but remain unknown to the customers. We model this dynamic by assuming the arrival rate Λ is drawn from a distribution, $\Lambda \sim \mathcal{A}$, where the server knows the realization of Λ, but the customer only knows the overall distribution $\mathcal{A}$.

1.1. Our findings

In this work, we study the problem of characterizing the structure of the optimal congestion signaling mechanism in a FIFO single-server queue. Customers arrive according to a Poisson process with rate Λ and have an exponential service requirement with mean 1. Here, the expected wait time is the same as the queue length. We consider the setting where the queue length (denoted by Queue) is unobservable to customers, all of whom are delay sensitive. For simplicity of illustration, in this section, we consider that the arrival rate Λ

is either λ_{High} or λ_{Low} with probability 0.5, and there are $K = 2$ types of customers, who receive utility upon service, given by $\text{Util}_i(\text{Queue}) = \text{Patience}_i - \text{Queue}$, for $i \in \{1, 2\}$. Here, $\text{Patience}_1 < \text{Patience}_2$, meaning that type-1 customers have lower patience than type-2 customers. Note that the expected utility is positive if and only if the expected wait time is below the customer's patience level.

The signaling mechanism is composed of two key components: (i) a set of signals $\mathcal{S}$, and (ii) a policy $\boldsymbol{\sigma}$ which maps the current system state, defined by the queue length and the arrival rate, to a signal, denoted by S. We follow the Bayesian persuasion framework to model the customer's response to the provided signal. We assume that service provider *commits* to a fixed signaling mechanism and the signaling mechanism is public knowledge. The Bayesian persuasion framework implies that the customers know the Bayes posterior of the expected utility from entering, i.e., $\mathbb{E}[\text{Util}_i(\text{Queue})|S]$, for every signal S within the set $\mathcal{S}$, and enter if the expected utility under the given signal is positive, i.e., if their conditional expected wait time is less than their patience.

Our first result, given in Proposition 1, provides a type of revelation principle (Krishna 2009), ensuring that it suffices for the server to specify to customers, given any system state, a *patience cutoff* such that only the customers with patience above the cutoff should enter. It guarantees that for any signaling mechanism and the corresponding customer equilibrium, there exists an equivalent signaling mechanism with $K+1 = 3$ congestion signals, each of which unambiguously conveys which customer types should enter (all types with patience above a certain cutoff), and following these recommendations is a customer equilibrium. For the three congestion signals, the "low congestion" signal implies an expected wait time (weakly) less than Patience_1, encouraging both customer types to enter. The "medium congestion" signal implies an expected wait time (weakly) between Patience_1 and Patience_2 and therefore dissuades type-1 customers from entering while encouraging type-2 customers to enter. Finally, the "high congestion" signal implies an expected wait time exceeds Patience_2 and thus suggests that both customer types do not enter.

Constrained-MDP (C-MDP): In Proposition 2, we present a C-MDP formulation of the problem. Given that the server uses patience cutoff signals (equivalent to the server's recommendation of the customer's action) and customers are obedient (following the server's recommendation is a customer equilibrium), we map the problem of finding the optimal signaling mechanism to solving for the optimal policy in a C-MDP. Specifically, we model the signals as the server's "action" and throughput as a reward function. The incentive compatibility of the customer equilibrium translates to constraints in the C-MDP. We exploit the C-MDP formulation as an analytical tool to develop a framework for determining the structure of the optimal policy.

Laminar structure: Theorem 1 characterizes the structure of the optimal signaling mechanism. We show that there exists an optimal policy $\boldsymbol{\sigma}$ such that, for each value of $\lambda \in \{\lambda_{\text{Low}}, \lambda_{\text{High}}\}$, the policy $\boldsymbol{\sigma}$ (given λ) has a *ordered laminar* structure.

We now informally introduce the laminar structure. For simplicity, we will ignore the discreteness of queue lengths in our description. (In our formal treatment, we handle discreteness by randomizing between consecutive integer queue lengths as necessary.) Under two customer types, for any given arrival rate $\lambda \in \{\lambda_{\text{Low}}, \lambda_{\text{High}}\}$, there exist ordered *laminar intervals*, Interval_1 (associated to "low congestion") and Interval_2 (associated to "medium congestion" signal), where ordered laminar means that the intervals are either non-overlapping ($\text{Interval}_1 \cap \text{Interval}_2 = \emptyset$) or Interval_1 contains Interval_2 ($\text{Interval}_2 \subsetneq \text{Interval}_1$). The signaling policy has an ordered laminar structure if the server provides the "medium congestion" signal when Queue is in Interval_2, and a "low congestion" signal if Queue is in Interval_1 but not in Interval_2. In some cases where Interval_1 strictly contains Interval_2, the policy has a counter-intuitive structure where the queue length interval for the "medium congestion" signal is *nested* within the interval for the "low congestion" signal.

Ordered laminar structure: Compared to the laminar structure seen in previous works, in this work, we demonstrate that the laminar intervals associated with different signals follow a strict nesting order based on their expected wait times. Specifically, a signal with a higher expected wait time ("medium congestion" signal) can have its corresponding interval nested within the interval of a signal with a lower expected wait time ("low congestion" signal), but not vice versa, i.e., we may observe $\text{Interval}_2 \subsetneq \text{Interval}_1$, but never $\text{Interval}_1 \subsetneq \text{Interval}_2$. A visual depiction of a laminar structure with nested intervals is provided in Figure 1.

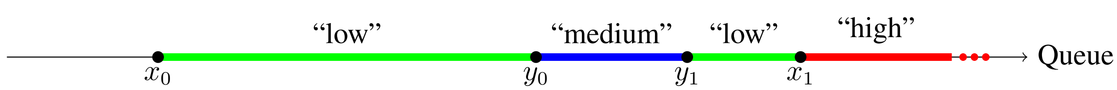

Figure 1 Pictorial representation of a (non-monotonic) laminar structure with $\text{Interval}_1 = [x_0, x_1]$ and $\text{Interval}_2 = [y_0, y_1]$.

When the policy exhibits a laminar structure, it can be fully characterized by specifying just two endpoints for each laminar interval (i.e., Interval_1 and Interval_2), which significantly reduces the complexity of the policy description. Moreover, as we demonstrate later, this laminar structure generalizes monotone policies, which are similar to threshold-based policies.

Intuition behind the non-monotonic laminar structure: The non-monotonic structure of the optimal policy emerges due to the misalignment between the server's and the customer's objectives. As the customers are only sensitive to the expected utility, the non-monotonic structure of the policy can still be compatible with their incentives. For example, consider the Bayes posterior of the expected utility under the "low congestion" signal in Figure 1. In this case, the negative expected utility (expected wait time above Patience_1) the (type-1) customers receive when they enter in the "bad region" (the interval (y_1, x_1)) is balanced by the positive expected utility (expected wait time below Patience_1) they receive in the "good region" (the interval (x_0, y_0)). This implies that a signaling policy with a non-monotonic structure can be compatible with the customer's incentives, hence, providing the server with more flexibility in designing the signaling policy.

For short queue lengths, the "low congestion" signal encourages more customers to join the system and boosts arrival rates. However, this quickly increases congestion, so at a certain threshold y_0, the server switches to signaling "medium congestion" to slow down the entry rate. However, to ensure that the expected wait time under the "medium congestion" signal is below Patience_2, the server is forced to stop providing "medium congestion" signal beyond the threshold y_1. As such, the server switches back to signaling "low congestion" for a small interval of queue length to extract more reward (in terms of throughput) up to the maximum threshold x_1 while ensure that the expected wait time under "low congestion" signal is below Patience_1. This results in a non-monotonic structure in the optimal policy.

Demand variation with homogeneous customers: The laminar structure of the optimal policy arises due to heterogeneous customers with private patience information. In the case of homogeneous customers (i.e., $\text{Patience}_1 = \text{Patience}_2$), the optimal signaling mechanism reduces to a binary signal (e.g., "low congestion" or "high congestion") and a threshold policy, where the threshold depends on the realization of Λ. Specifically, for λ_{Low} and λ_{High}, there exist thresholds T_{Low} and T_{High} respectively, such that when the arrival rate realization is λ_{Low} (or λ_{High}), the server signals the customer to "Enter" (same as "low congestion" signal) if Queue $\leq T_{\text{Low}}$ (or Queue $\leq T_{\text{High}}$), and signals the customer to "Leave" (same as "high congestion" signal) otherwise. These results align with those in Lingenbrink and Iyer (2019), albeit without considering demand variation.

Further, as the customers are unaware of the realization of the arrival rate Λ, the server can leverage this information gap to achieve higher throughput, as compared to the case when customers know the realized value of Λ. Intuitively, this occurs because randomness in Λ enables the server to adopt a signaling mechanism with more flexibility in meeting the incentive compatibility constraints. For example, the server may choose the thresholds T_{Low} and T_{High} such that the customers receive a strictly positive expected utility from entry on some days (i.e., $\mathbb{E}[\text{Utility}|\text{“Enter”}, \Lambda = \lambda_{\text{High}}] > 0$) and compensated by providing negative expected utility from entry on the other days (i.e., $\mathbb{E}[\text{Utility}|\text{“Enter”}, \Lambda = \lambda_{\text{Low}}] < 0$), while maintaining $\mathbb{E}[\text{Utility}|\text{“Enter”}] \geq 0$. As presented in Proposition 5, this flexibility ensures that the server achieves higher throughput when the customers do not know the arrival rate realization.

Interestingly, we also identify examples where the server might benefit from randomness in Λ, in the sense that the maximum achievable throughput when $\Lambda \sim \text{Unif}\{\lambda_{\text{Low}}, \lambda_{\text{Higher}}\}$ is larger than that in the scenario where the arrival rate is fixed to the expected value $\bar{\lambda} = \frac{1}{2}(\lambda_{\text{Low}} + \lambda_{\text{High}})$. However, this phenomenon typically occurs only when the arrival rate variation is small, i.e., λ_{Low} is close to λ_{High}, as large fluctuations lead to substantial throughput losses on days with low demand.

1.2. Literature review

The problem of signaling between a sender and a receiver dates back to the classic cheap talk game introduced by Crawford and Sobel (1982), where the authors analyze the communication between a sender (who has information) and a receiver (who takes action based on the information). The key aspect of this interaction is that the sender and receiver have partially aligned interests, meaning that while they both benefit from the receiver taking a more informed action, their preferences are not completely aligned. The paper shows how the informativeness of communication depends on the misalignment between the sender's and receiver's interests, leading to equilibrium outcomes that range from fully informative to completely uninformative. Since the work by Crawford and Sobel (1982), there have been several works on the topic of strategic communication and we refer the readers to Farrell and Rabin (1996), Kartik (2009), Krishna and Morgan (2001), and the references therein.

Building on this foundation, Allon et al. (2011) investigate the impact of strategic delay signals on customer behavior and operational efficiency in service systems, where each customer's actions have externalities on system dynamics, implying that the prior distribution of system's state is not exogenous (as in Crawford and Sobel (1982)), but is determined endogenously. They analyze the equilibrium behavior of the server and the customer and argue that information provisioning (providing incomplete information) improves the firm's profits and the customers' expected utility.

In subsequent work, and closest to our work, Lingenbrink and Iyer (2019) consider the problem of signaling mechanism design in a single-server queue with homogeneous customer types and a revenue-maximizing server. They consider the equilibrium behavior of the system (similar to Allon et al. (2011)) and use the Bayesian persuasion framework to model the customer's response, which reduces the signaling mechanism design problem to a linear optimization problem over the resulting stationary distribution. They show that a binary signaling mechanism with a threshold structure is optimal. Further extensions include Anunrojwong et al. (2024), which assumes that customers are risk-sensitive, and Anunrojwong et al. (2023), focusing on welfare maximization.

Similar to Allon et al. (2011) and Lingenbrink and Iyer (2019), we consider a single-server queue, where the state distribution (i.e., steady-state distribution in our setting) depends endogenously on the chosen signaling mechanism. However, our work also captures additional real-world features, such as demand variation and customer heterogeneity with private types. We show that with these features, the optimal signaling mechanism can have a more complicated (laminar) structure than the simple threshold structure established in Lingenbrink and Iyer (2019).

Concerning the policy structure seen in this work, Candogan and Strack (2023) and Candogan (2020) also obtain a laminar structure for the optimal signaling mechanism for a system with heterogeneous customers and linear rewards. However, similar to that in Crawford and Sobel (1982), they consider a fixed prior, which is independent of the signaling mechanism. In this work, the prior distribution is endogenous, which requires us to employ a different technique than those used in Candogan and Strack (2023) and Candogan (2020). In particular, we reformulate the problem as a C-MDP and study its properties to derive the structure of the optimal policy. This analysis tool can be of independent interest. Moreover, we establish a strict ordering on how the laminar intervals corresponding to different signals nest within each other.

Apart from the literature mentioned above, our work also relates to the following sets of literature.

Delay announcement in service systems: The question of whether to reveal queue length information to customers has been a significant topic of research for many service systems, such as call centers (Akşin et al. (2017), Yu et al. (2017)), hospital emergency departments (Dong et al. (2019)), and theme parks (Kostami and Ward (2009), Nirenberg et al. (2018)). One of the early works examining this issue is Hassin (1986), where the authors show that it might benefit a revenue-maximizing server to disclose the queue length information to an incoming customer. Whitt (1999) studies the setting of call centers and shows that the average wait time can be reduced when accurate delay information is provided. Guo and Zipkin (2007) considers three levels of information and shows that more information may or may not be beneficial, depending on the distribution of customer's sensitivity to delay.

Several works consider the effect of delay announcement on customer satisfaction and wait time with features such as customer abandonment (Armony et al. (2009), Akşin et al. (2017)) and a call back option (Armony and Maglaras 2004). The literature on delay announcements is vast, and we refer the readers to the survey Ibrahim (2018) for more details.

Strategic behavior in queues: There is an increasing amount of literature on the strategic behavior of agents in a queueing system. In the seminal work, Naor (1969) considers the problem of maximizing revenue or welfare in the tolling system, which is modeled as a single server queue ($M/M/1$ queue) with *observable* queue length. Edelson and Hilderbrand (1975) consider the same problem as in Naor (1969), but in a setting where the queues are unobservable. Chen and Frank (2001) extend the classic model of Naor (1969) and show that even though when the customers are homogeneous, the revenue optimizing prices also maximizes social welfare, the same does not hold when the customers are heterogeneous. The book Hassin and Haviv (2003) provides a nice overview of the classic results and methodology in this area. For more details, the readers can also refer to Hassin (2016) and Economou (2021).

Bayesian Persuasion: In this work, we use the Bayesian persuasion framework to model the effect of the signaling mechanism on customers' behavior. Bayesian persuasion is a theoretical framework that explores how an informed sender can strategically influence the actions of a less informed receiver by controlling the information. Introduced in the seminal work Kamenica and Gentzkow (2011), this framework models scenarios where the sender commits to a signaling scheme to affect the receiver's beliefs and subsequent actions. Kolotilin (2018) examines the problem of optimal information disclosure by providing a linear programming framework to determine the optimal policy. Dworczak and Martini (2019) presents a price-theoretic model for Bayesian persuasion and characterizes the optimal information structure. Kolotilin et al. (2017), Candogan and Strack (2023) considers the scenario where receivers have private type information that cannot be observed by the sender. Bergemann and Morris (2019) presents a unified framework connecting the models in Bayesian persuasion, signaling games, and mechanism design. See Rayo and Segal (2010), Bergemann and Morris (2016a,b) and Arieli and Babichenko (2019) for more details.

Constrained-MDP: In this work, we use the theory of C-MDP to derive the structural properties of the optimal signaling mechanism. Altman (2021) provides a comprehensive overview of C-MDP, and several arguments in this work are based on the results presented in Altman (2021). There is an extensive body of literature developing algorithms to learn the optimal policy of constrained MDP (see, e.g., Efroni et al. (2020), Wei et al. (2024), Chen et al. (2021)). In this work, we use a Lagrangian approach that maps the C-MDP problem to an admission control problem using strong duality. Khan and Subramanian (2023) and Altman (2021) present a strong duality argument for a class of C-MDPs and constrained Partially Observable MDPs, and some of the arguments to establish strong duality (see Lemma 2) in this work are inspired by the arguments presented in Khan and Subramanian (2023).

1.3. Paper outline

We begin by presenting the model in Section 2. The main results are outlined in Section 3, which includes three key components: (i) Proposition 1 identifies the set of signals used by the server, (ii) Proposition 2 introduces the C-MDP formulation for the throughput maximization problem, and (iii) Theorem 1 describes the structure of the optimal policy. In Sections 4 and 5, we explore the individual effects of customer heterogeneity and demand variation, respectively. Section 6 provides the mathematical details of Theorem 1, where the C-MDP formulation is used to characterize the optimal policy. Finally, in Section 7, we conclude with additional discussion and propose directions for future research.

1.4. Basic notations

We use $\mathbb{1}_d$ to denote a vector of all ones in $\mathbb{R}^d$ and, $\mathbb{1}_\infty$ for the same in $\mathbb{R}^\infty$. $\mathbb{Z}_+$ denotes the set of non-negative integers. Bold letter denote vectors in $\mathbb{R}^\infty$, e.g., $\mathbf{x}$ denote the vector $\{x_n\}_{n=0}^\infty$. And for any vector $\mathbf{x}$, x_n denotes the n^{th} element of the vector. The operation $\langle\cdot,\cdot\rangle$ denote the dot product of two vectors in $\mathbb{R}^\infty$, i.e., $\langle \mathbf{x}, \mathbf{y}\rangle = \sum_{n=0}^\infty x_n y_n$. We use $\mathbf{I}$ to denote the identity function, i.e., $\mathbf{I}(n) = n$, and for any vector $\mathbf{x} \in \mathbb{R}^\infty$, we use $\langle \mathbf{x}, \mathbb{1}\rangle = \sum_{n=1}^\infty n x_n$. For any positive integer d, we use $[d]$ to denote the set $\{1, 2, \ldots, d\}$. For any set $\mathcal{T} \subset \mathbb{R}$, $|\mathcal{T}|$ denotes the Lebesgue measure of the set. Unless stated otherwise, X is a uniform random variable in $[0, 1]$, i.e., $X \sim \text{Uniform}[0, 1]$.

2. Model

We consider the setting of a single-server queue, where customers arrive according to a Poisson process with rate Λ and have exponential service requirement with mean 1. The arrival rate Λ is a random variable and follows the distribution $\mathcal{A}$, assumed to have finite support $\mathcal{L} = \{\lambda_1, \lambda_2, \ldots, \lambda_L\}$, where $L < \infty$, and $\mathbb{P}(\Lambda = \lambda_l) = q_l$.

The sequence of events is as follows: First, an arrival rate is drawn from the distribution $\mathcal{A}$, and the server knows the realized value of Λ, say $\Lambda = \lambda$. Customers then arrive at the system at rate λ and must decide whether to enter or not. However, they cannot observe the system's current state, which includes both the current queue length, denoted by W, and the arrival rate λ. Instead, the server looks at the system's current state and provides coarse congestion information, referred to as a *signal*, using a publicly known *signaling mechanism*. Based on the provided signal and the knowledge of the system – the signaling mechanism and the system parameters – and anticipated strategy of other customers, the customers form a *belief* about their expected utility (same as the Bayes posterior of the expected utility) and choose to enter the system if they expect a positive utility; otherwise, they decide to leave (not enter).

Signaling mechanism: The service provider commits to a signaling mechanism, denoted by $\mathbf{\Sigma}$, which sends "informative" signals about the system's state. In particular, the signaling mechanism $\mathbf{\Sigma} = (\mathcal{S}, \boldsymbol{\sigma})$ consists of a set of signals $\mathcal{S}$ and a policy $\boldsymbol{\sigma} : \mathbb{Z}_+ \times \mathcal{L} \to \Omega(\mathcal{S})$, where $\Omega(\mathcal{S})$ denotes the set of probability distributions over the set of signals $\mathcal{S}$. Mathematically, if the state is $(W, \Lambda) = (n, \lambda)$,

$$\sigma(S|n, \lambda) := \mathbb{P}\big(\xi(S) = 1 | W = n, \Lambda = \lambda\big),$$

where $\xi(S)$ is an indicator denoting the event that the server sends the signal S. Without loss of generality, we denote the signaling mechanism as $\mathbf{\Sigma} = \{\mathbf{\Sigma}^1, \mathbf{\Sigma}^2, \ldots, \mathbf{\Sigma}^L\} = \{(\mathcal{S}, \boldsymbol{\sigma}^l)\}_{l \in [L]}$, such that if the realized arrival rate is λ_l, the server uses the corresponding mechanism $\mathbf{\Sigma}^l = (\mathcal{S}, \boldsymbol{\sigma}^l)$, where $\sigma^l(S|n) := \sigma(S|n, \lambda_l)$.

Customer's utility: We assume that the customers are heterogeneous and differ in terms of the utility they receive upon service. Each arriving customer belongs to one of the K types, where customer is of type i with probability p_i for $i \in [K]$. For ease of notation, we define $P_i := \sum_{j=i}^{K} p_j$ and $P_{K+1} = 0$. If a type-i customer chooses to join the queue when its length is W, their expected utility from receiving service is $h_i(W)$. However, if the customer decides not to join the queue, their utility is zero. The utility function $h_i(W)$ is linear in W, and is given by

$$h_i(W) := c_i - b_i W, \ \forall i \in [K],$$

where c_i's and b_i's are constants, and the ratio c_i/b_i denotes the *patience* of customer type i, meaning that the customers receive positive utility if the wait time is below their patience level, and vice versa. Linear utility functions are commonly used in the queueing literature (see, e.g., Allon et al. 2011, Naor 1969). Without loss of generality, we assume that $0 < \frac{c_1}{b_1} < \frac{c_2}{b_2} < \cdots < \frac{c_K}{b_K}$, i.e., type-1 is the least patient type and type K is the most patient type.

Customer strategy: In this work, the customers are strategic and Bayesian, who aim to maximize the expected utility. Upon arrival, the customers have two actions available: enter or leave. We assume that there is a single decision epoch for the customers, that is, once they enter, they remain in the queue until they receive service. In this case, given a server's signaling mechanism $\mathbf{\Sigma}$, a *customer strategy* is a mixed strategy that maps a signal to a probability distribution over the action: "Enter" (1) or "Leave" (0). A customer strategy vector is denoted by a mapping $\boldsymbol{\nu} : \mathcal{S} \to [0,1]^K$, where $\nu_i(S)$ is the probability with which a type-i customer enters the queue given the signal S. (We will assume that all type i customers adopt the same, possibly mixed, strategy.) Note that the strategy vector $\boldsymbol{\nu}$ is a pure strategy if $\boldsymbol{\nu} : \mathcal{S} \to \{0,1\}^K$.

Stationary distribution given the customer strategy $\mathbf{\Pi}(\mathbf{\Sigma},\boldsymbol{\nu})$: Under a given signaling mechanism $\mathbf{\Sigma}$ and a customer strategy vector $\boldsymbol{\nu}$, the system dynamics can be described as a birth-death Markov chain. In this work, we assume that the server's signaling mechanism $\mathbf{\Sigma}$ and the customer strategy vector $\boldsymbol{\nu}$ are such that the steady-state queue length distribution exists for all realizations of Λ. We use $\mathbf{\Pi}(\mathbf{\Sigma},\boldsymbol{\nu}) = \{\boldsymbol{\pi}^l\}_{l\in[L]}$ to denote the tuple of steady-state distributions, where $\boldsymbol{\pi}^l = \{\pi_n^l\}_{n=0}^{\infty}$ denotes the steady-state distribution of the queue length given the arrival rate realization is $\Lambda = \lambda_l$. For each $l \in [L]$, $\boldsymbol{\pi}^l$ is the solution of the following flow balance equation.

$$\pi_{n+1}^l = \lambda_l \pi_n^l \sum_{S\in\mathcal{S}} \sigma^l(S|n) \sum_{i=1}^{K} p_i \nu_i(S),\ \forall n \geq 0, \ \text{ and } \ \sum_{n=0}^{\infty} \pi_n^l = 1. \tag{1}$$

When the context is clear, we typically omit the dependency of the tuple of steady-state distributions on $\mathbf{\Sigma}$ and $\boldsymbol{\nu}$ and use the notation $\mathbf{\Pi}$ for convenience.

Customer equilibrium: In this work, similar to Lingenbrink and Iyer (2019), we introduce a notion of customer equilibrium, and we focus on the setting of symmetric customer equilibrium, meaning an equilibrium where each type-i customer uses the same mixed strategy $\nu_i(S)$ upon receiving signal S. An important implication of this assumption is that, for any customer type i, for the signals under which type i is indifferent between entering or leaving, all customers of type i use the same tie-breaking rule.

As the customers arrive according to a Poisson process, using the PASTA property (Wolff 1982), an incoming customer sees a queue length drawn from the steady-state distribution. As such, the customer's *prior* belief of the system's state is given by $(\Lambda, W) \sim (\mathcal{A}, \mathbf{\Pi})$, that is, the prior is given by $\mathbb{P}(\Lambda = \lambda_l, W = n) = q_l \pi_n^l$ for all $l \in [L]$ and $n \geq 0$. Upon receiving the signal S, the customer uses Bayes' rule to form a posterior belief of the expected utility, given by $\mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}[h_i(W)|\xi(S) = 1]$, where $\mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}[\cdot]$ denotes the expectation under the distribution $(\mathcal{A}, \mathbf{\Pi})$, and the policy $\boldsymbol{\sigma}$. Mathematically, for any function $f(W, \Lambda)$, we have

$$\mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}[f(W,\Lambda)] = \sum_{n=0}^{\infty} \sum_{l\in[L]} q_l \pi_n^l f(n, \lambda_l), \ \text{ and}$$

$$\mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}[f(W,\Lambda)|\xi(S) = 1] = \frac{\sum_{n=0}^{\infty} \sum_{l\in[L]} q_l \pi_n^l \sigma^l(S|n) f(n, \lambda_l)}{\sum_{n=0}^{\infty} \sum_{l\in[L]} q_l \pi_n^l \sigma^l(S|n)} \ \text{ whenever } \sum_{n=0}^{\infty} \sum_{l\in[L]} q_l \pi_n^l \sigma^l(S|n) > 0.$$

Definition 1 (Customer equilibrium)**.** *Given a server's signaling mechanism $\mathbf{\Sigma}$, the customer strategy vector $\boldsymbol{\nu}$ is a customer equilibrium if, for each customer type i, the strategy $\nu_i(\cdot)$ is a best response to $\mathbf{\Sigma}$ and the tuple of steady-state distributions $\mathbf{\Pi} = \mathbf{\Pi}(\mathbf{\Sigma},\boldsymbol{\nu})$, i.e., for all $i \in [K]$ and $S \in \mathcal{S}$, we have*

$$\nu_i(S) = \begin{cases} 1 & \text{if } \mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}[\xi(S)h_i(W)] > 0, \\ 0 & \text{if } \mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}[\xi(S)h_i(W)] < 0, \end{cases} \tag{2}$$

and $\nu_i(S) \in [0,1]$ if $\mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}[\xi(S)h_i(W)] = 0$.

Note that under any given signal S, a customer will choose to enter if they receive positive expected utility conditioned on receiving signal S. However, in Eq. (2) of Definition 1, we avoid using conditional expectations. This ensures that we have a well-defined customer equilibrium even for signals that are sent with zero probability. We say that a signal $S \in \mathcal{S}$ is *present* (given $\boldsymbol{\Sigma}$, $\boldsymbol{\nu}$ and $\boldsymbol{\Pi}(\boldsymbol{\Sigma}, \boldsymbol{\nu})$) if there is a positive probability of giving signal S, i.e., $\mathbb{E}_{\boldsymbol{\Pi},\boldsymbol{\Sigma}}\big[\xi(S)\big] = \sum_{n=0}^{\infty} \sum_{l \in [L]} q_l \pi_n^l \sigma^l(S|n) > 0$, and *not present* otherwise. For any signal that is present, the condition $\mathbb{E}_{\boldsymbol{\Pi},\boldsymbol{\Sigma}}[\xi(S) h_i(W)] > 0$ is equivalent to $\mathbb{E}_{\boldsymbol{\Pi},\boldsymbol{\Sigma}}[h_i(W)|\xi(S) = 1] > 0$. Similar argument holds when $\mathbb{E}_{\boldsymbol{\Pi},\boldsymbol{\Sigma}}[\xi(S) h_i(W)] < 0$. For any signal that is not present, the condition $\mathbb{E}_{\boldsymbol{\Pi},\boldsymbol{\Sigma}}[\xi(S) h_i(W)] = 0$ is trivially satisfied, which is consistent with the fact that any mixed strategy $\nu_i(S) \in [0,1]$ is a best response if the signal S is not present.

In this work, we assume the existence of a customer equilibrium that satisfies Eq. (2). Formally proving this requires demonstrating the existence of a fixed point for an operator implicitly defined by Eq. (2) alongside the flow balance equations in Eq. (1). For a detailed argument on the existence of such an equilibrium, we refer readers to the results in Hopenhayn (1992) and Hopenhayn and Rogerson (1993), which are applicable in the context of this work.

Server's objective: For a given signaling mechanism $\boldsymbol{\Sigma}$ and a customer strategy vector $\boldsymbol{\nu}$, the steady-state throughput is given by

$$\text{Th}(\boldsymbol{\Sigma}, \boldsymbol{\nu}) = \sum_{l=1}^{L} q_l \Big[\sum_{n=0}^{\infty} \pi_n^l \sum_{S \in \mathcal{S}} \lambda_l \sigma^l(S|n) \sum_{i=1}^{K} p_i \nu_i(S) \Big],$$

where $\{\boldsymbol{\pi}^l\}_{l \in [L]}$ is the solution of the flow balance equations in Eq. (1). The service provider's goal is to design a signaling mechanism $\boldsymbol{\Sigma}$ that results in a customer equilibrium $\boldsymbol{\nu}$, which maximizes the system throughput. Note that there may be multiple customer equilibria. In such a situation, we assume the server can choose a "focal" customer equilibrium (among all possible equilibria), motivated by the "revelation principle" (Proposition 1) which we establish in the next section.

3. Main results

We now provide the main results of this work. Section 3.1 provides a revelation principle, which tells us that the server needs to deploy only $K+1$ distinct signals, each of which implicitly recommends a specific action (enter or not) to each customer type. Section 3.2 provides the C-MDP formulation of the server's throughput maximization problem, and finally, Section 3.3 presents the structure of the policy under the optimal signaling mechanism.

3.1. Preliminary observations

We first introduce three definitions which will enable us to state our revelation principle.

As the customer type is private information, the server designs a signaling mechanism $\boldsymbol{\Sigma}$ that provides the same signal to all the customers, independent of the type of the arriving customer. This ensures that an incoming customer receives a signal that depends only on the system's state (W, Λ). The server has access to possibly infinite number of signals, so the optimal signaling mechanism may not be unique. Nonetheless, the customer's response to any given signal is limited to two choices: "Enter" or "Leave". This allows us to narrow down the space of signaling mechanisms and the corresponding customer equilibrium. To achieve this, we define an equivalence relation between two signaling mechanisms.

Definition 2 (Equivalence of signaling mechanisms)**.** *Two signaling mechanisms,* $\mathbf{\Sigma} = (\mathcal{S}, \boldsymbol{\sigma})$ *and* $\tilde{\mathbf{\Sigma}} = (\tilde{\mathcal{S}}, \tilde{\boldsymbol{\sigma}})$*, and the respective customer equilibria* $\boldsymbol{\nu}$ *and* $\tilde{\boldsymbol{\nu}}$*, are said to be* equivalent *if their respective steady-state distribution tuples* $\mathbf{\Pi}$ *and* $\tilde{\mathbf{\Pi}}$ *are equal, i.e.,* $\boldsymbol{\pi}^l = \tilde{\boldsymbol{\pi}}^l$ *for all* $l \in [L]$*.*

Definition 2 is the same as the definition of equivalence of two signaling mechanisms provided in Lingenbrink and Iyer (2019). A similar definition is also used in (Allon et al. 2011, Definition 3.3). A notion of equivalence is required because two different signals (e.g., phrases) may imply the same expected posterior utility. Using the equivalence relation in Definition 2, we now bound the number of signals needed by the server, and moreover show that it suffices to use signals with a simple interpretation in terms of recommended customer actions.

Definition 3 (Patience cutoff signals)**.** *Given* K *customer types with distinct patience levels* $\{c_i/b_i\}_{i\in[K]}$*, the* patience cutoff signals *(or recommendations) are the set of signals*

$$\mathcal{S}^{[K]} = \{S_1, S_2, \ldots, S_{K+1}\} := \Big\{S = (s_1, s_2, \ldots, s_K) \in \{0,1\}^K : s_1 \le s_2 \le \cdots \le s_K\Big\}, \tag{3}$$

where $S_i := \{s_1 = 0, \ldots, s_{i-1} = 0, s_i = 1, \ldots, s_K = 1\}$ for $i \in [K]$*, and* $S_{K+1} = \{0, 0, \ldots, 0\}$*. Here, the* k^{th} *element* s_k *is interpreted as the server's binary action recommendation to type-*k *customers (a 1 says "Enter" and a 0 says "Leave"). Hence the signal* S_i *is interpreted as a recommendation for customer types* $1, 2, \ldots, i-1$ *to leave, and types* $i, \ldots, K$ *to enter (i.e., only customers with patience level larger than or equal to* c_i/b_i *are recommended to enter).*

To illustrate, consider $K = 2$. The set of patience cutoff signals is $\mathcal{S}^{[2]} = \big\{S_1 := \{1,1\}, S_2 := \{0,1\}, S_3 := \{0,0\}\big\}$, where S_1 means "both types enter", S_2 means "only type-2 enters", and S_3 means "no one enters".

Definition 4 (Obedience)**.** *Given that the set of signals is the patience cutoff signals* $\mathcal{S}^{[K]}$*, we define the* obedient *customer strategy vector (in short,* obedience*) as the one in which, given signal* $S_i = \{s_1, s_2, \ldots, s_K\}$*, each type-*$k$ *customer follows its patience cutoff recommendation* s_k *to enter or not. Formally, the obedient customer strategies are, for any signal* $S_i \in \mathcal{S}^{[K]}$ *and* $k \in [K]$*,* $\nu_k(S_i) = 1$ *if* $k \ge i$ *and* $\nu_k(S_i) = 0$ *if* $k < i$*.*

Note that in the design of patience cutoff signals in Definition 3 and obedient customer strategy in Definition 4, we are implicitly using that the customer patience levels c_i/b_i's follow a strictly ascending order. More specifically, we are using the following: Given any signaling mechanism $\mathbf{\Sigma} = (\mathcal{S}, \boldsymbol{\sigma})$, in a customer equilibrium, for any signal S, if a type-i customer's strategy is to enter the system upon receiving signal S, then for all $k \ge i$, a type-k customer's best response is also to enter under signal S. This gives the condition $\nu_1(S) \le \nu_2(S) \le \cdots \le \nu_K(S)$, meaning that any customer equilibrium with pure strategies satisfies $\boldsymbol{\nu}(S) \in \mathcal{S}^{[K]}$ for any $S \in \mathcal{S}$, providing an intuitive reasoning behind choosing $\mathcal{S}^{[K]}$ itself as the set of signals.

With these definitions in place, we are now ready to state our revelation principle.

Proposition 1 (Revelation principle with $K+1$ signals)**.** *For any signaling mechanism $\mathbf{\Sigma}' = (\mathcal{S}, \boldsymbol{\sigma}')$ and a customer equilibrium $\boldsymbol{\nu}'$, there exists an equivalent* recommendation-obedient *signaling mechanism, i.e., a signaling mechanism $\mathbf{\Sigma} = (\mathcal{S}^{[K]}, \boldsymbol{\sigma})$ which uses the patience cutoff signals $S^{[K]}$ such that obedience is a customer equilibrium under $\mathbf{\Sigma}$.*

Proposition 1 establishes a framework for limiting the set of signals the server uses and relates to the *revelation principle* (Krishna 2009) from mechanism design. In our context, it states that it suffices for the server to use the patience cutoff signals $\mathcal{S}^{[K]}$ and design its policy $\mathbf{\Sigma}$ such that following the server's recommendations is a customer equilibrium. To ensure customer obedience, the server must design the policy $\boldsymbol{\sigma}$ to align with the customers' incentives. This requirement is captured by the constraints of the C-MDP formulation, which we introduce in Section 3.2. Notably, the result in Proposition 1 does not depend on the server's objective.

With a single customer type ($K=1$) and no demand variation ($L=1$), Proposition 1 was proved in Lingenbrink and Iyer (2019). We extend their argument to a setting with heterogeneous customers with private patience information and demand variation. The complete proof of Proposition 1 is provided in Appendix C.1 and we provide a sketch of the proof strategy next.

First, we show that, for any signaling mechanism and a corresponding customer equilibrium, the server can always design an equivalent signaling mechanism where customer equilibrium is a pure strategy, i.e., $\boldsymbol{\nu}: \mathcal{S} \to \{0,1\}^K$. Intuitively, the customer-side randomization can be integrated into the server's signaling process, while still yielding the same stationary distribution. Further, using the fact that c_i/b_i follow an ascending order, we have $\boldsymbol{\nu}: \mathcal{S} \to \mathcal{S}^{[K]}$. Second, we note that if a signaling mechanism $\mathbf{\Sigma}'$ consists of two signals, S and S', so that for each customer type, the customer's response to both signals is the same, then the server can effectively combine the two signals to create a new signaling mechanism $\mathbf{\Sigma}$, equivalent to $\mathbf{\Sigma}'$. Thus, it suffices to design a signaling mechanism where each signal uniquely corresponds to a customer response vector in $\mathcal{S}^{[K]}$. Subsequently, we can relabel the signals such that the set of signals is $\mathcal{S}^{[K]}$ itself, and $\boldsymbol{\nu}(S) = S$ for all $S \in \mathcal{S}^{[K]}$.

3.2. Constrained MDP formulation of the problem

We now formulate the server's throughput maximization problem presented in Section 2 as a C-MDP. From here onwards, we leverage Proposition 1 to restrict consideration to recommendation-obedient signaling mechanisms $\mathbf{\Sigma} = (\mathcal{S}^{[K]}, \boldsymbol{\sigma})$, i.e., the signaling mechanism uses the patience cutoff signals and obedience is a customer equilibrium. We also replace the notation $\mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}[\cdot]$ by $\mathbb{E}_{\boldsymbol{\sigma}}[\cdot]$ for simplicity.

Proposition 2. *Fix the arrival rate distribution $\mathcal{A}$, the K customer patience levels $\{c_i/b_i\}_{i\in[K]}$ and the distribution $\{p_i\}_{i\in[K]}$ over them. The problem of finding the throughput-maximizing signaling mechanism* *reduces to the following reward maximization problem over signaling policies $\boldsymbol{\sigma}$, that operate with the patience cutoff signals $\mathcal{S}^{[K]}$:*

$$\begin{aligned} \textit{Throughput C-MDP:} \quad \max_{\boldsymbol{\sigma}} \quad & \mathbb{E}_{\boldsymbol{\sigma}}\Big[\sum_{i=1}^{K} \Lambda P_i \xi(S_i)\Big] \\ \textit{s.t.} \quad & \mathbb{E}_{\boldsymbol{\sigma}}\big[\xi(S_i) h_i(W)\big] \geq 0, \quad \forall i \in [K] \\ & \mathbb{E}_{\boldsymbol{\sigma}}\big[\xi(S_{i+1}) h_i(W)\big] \leq 0, \quad \forall i \in [K]. \end{aligned} \tag{4}$$

Proposition 2 uses the result in Proposition 1 and reduces the server's throughput maximization problem to finding an optimal policy $\boldsymbol{\sigma}$ for a C-MDP. In the C-MDP (4), the server chooses from the $K+1$ actions (signals) in the action space $\mathcal{S}^{[K]}$, and $\boldsymbol{\sigma} = \{\boldsymbol{\sigma}^l\}_{l\in[L]}$ denotes the policy. The C-MDP (4) includes the $2K$ *obedience* constraints, also known as *incentive compatibility* (IC) constraints, two for each of the K customer types, which ensures that obedience is a customer equilibrium (i.e., $\boldsymbol{\nu}(S) = S$ for all $S \in \mathcal{S}^{[K]}$).

The idea behind the IC constraints is as follows: As the server is using the patience cutoff signals, to ensure that obedience is a customer equilibrium, for any signal $S_i \in \mathcal{S}^{[K]}$, we need two things. First, under signal S_i, any customer with a patience level equal to or greater than that of the type-i customer should receive a non-negative expected utility if entering, giving us the constraint $\mathbb{E}_{\boldsymbol{\sigma}}\big[\xi(S_i)h_i(W)\big] \geq 0$. And second, any customer with a patience level equal to or less than that of type-$(i-1)$ customer should receive a non-positive expected utility if entering, giving us the constraint $\mathbb{E}_{\boldsymbol{\sigma}}\big[\xi(S_i)h_{i-1}(W)\big] \leq 0$. As these constraints must hold for all $S_i \in \mathcal{S}^{[K]}$, we get the IC constraints in (4).

3.3. Structure of the optimal signaling mechanism

We now present our main result on the structure of the optimal policy. We start by introducing the notion of representing the policy (which is defined on non-negative integers) using a partition on the real line. For any $l \in [L]$, policy $\boldsymbol{\sigma}^l$ can be represented using a partition $\{\mathcal{T}_i\}_{i\in[K+1]}$ of $\mathbb{R}_+$, such that $\sigma^l(S_i|n) = \mathbb{P}(n+X \in \mathcal{T}_i) = |\mathcal{T}_i \cap [n, n+1)|$, where $X \sim$ Uniform$[0,1]$. In such a case, the policy is called a *partitional policy* with respected to partition $\{\mathcal{T}_i\}_{i\in[K+1]}$. Intuitively, just as a probability distribution maps to a partition of $[0,1]$, a policy over non-negative integers maps to a partition of the real line. Mathematically, one possible representation is

$$\mathcal{T}_i = \bigcup_{n=0}^{\infty} \Big[n + \sum_{j=1}^{i-1} \sigma^l(S_j|n), n + \sum_{j=1}^{i} \sigma^l(S_j|n)\Big), \; \forall i \in [K+1]. \tag{5}$$

Given $\Lambda = \lambda_l$, upon a customer's arrival, if the queue length is n, the server generates $X \sim$ Uniform$[0,1]$ and then maps the value $n+X$ to a signal in $\mathcal{S}^{[K]}$, that is, it provides the signal S_i if $n+X \in \mathcal{T}_i$. Even though the policy $\boldsymbol{\sigma}^l$ can be represented using a deterministic partition $\{\mathcal{T}_i\}_{i\in[K+1]}$, the representation in Eq. (5) is highly complex. In this work, we show that the optimal policy of Throughput C-MDP (4) can be represented using an ordered laminar partition.

Definition 5 (Ordered Laminar Partition of $\mathbb{R}_+$)**.** *A set of intervals* $\{\mathcal{I}_1, \dots \mathcal{I}_{K+1}\}$*, where* $\mathcal{I}_i \subseteq \mathbb{R}_+$ *and* $\cup_{i \in [K+1]} \mathcal{I}_i = \mathbb{R}_+$*, is called an* ordered laminar interval family *if, for any pair of subsets* $\mathcal{I}_i$ *and* $\mathcal{I}_j$ *(*$i < j$*), either* $\mathcal{I}_i \cap \mathcal{I}_j = \emptyset$ *or* $\mathcal{I}_j \subseteq \mathcal{I}_i$*. Further, a partition* $\{\mathcal{T}_1, \dots \mathcal{T}_{K+1}\}$ *of* $\mathbb{R}_+$ *is called an ordered laminar partition of* $\mathbb{R}_+$ *if there exists an ordered laminar interval family* $\{\mathcal{I}_1, \dots \mathcal{I}_{K+1}\}$ *such that* $\mathcal{T}_i = \mathcal{I}_i \setminus \cup_{j>i} \mathcal{I}_j$*.*

The notation $\mathcal{T}_i = \mathcal{I}_i \setminus \cup_{j>i} \mathcal{I}_j$ implies that $\mathcal{T}_i$ is derived by removing, from $\mathcal{I}_i$, the intervals $\mathcal{I}_j$'s (with $j \neq i$) that are contained in $\mathcal{I}_i$. To illustrate, Figure 2 provides a counter-example and an example of a laminar partition. Importantly, our definition imposes a specific ordering on the intervals $\{\mathcal{I}_1, \dots \mathcal{I}_{K+1}\}$, i.e., we require that $\mathcal{I}_i \subseteq \mathcal{I}_j$ can occur only if $i < j$. A more general definition (that doesn't consider this ordering) is provided in Candogan and Strack (2023), from which the pictorial examples in Figure 2 are also borrowed.

Figure 2 The partition on the right is a laminar partition corresponding to the laminar interval family $\mathcal{I}_1 = [0, \infty)$, $\mathcal{I}_2 = [1, 2)$ and $\mathcal{I}_3 = [3, 4)$, while the partition on the left is not a laminar partition.

Definition 6 (Ordered Laminar Policy)**.** *For any given* $l \in [L]$*, a policy* $\boldsymbol{\sigma}^l$ *(that uses patience cutoff signals) is* ordered laminar *(or has an ordered laminar structure) if there exist an ordered laminar partition* $\{\mathcal{T}_1, \dots \mathcal{T}_{K+1}\}$ *such that* $\sigma^l(S_i|n) = \mathbb{P}(n + X \in \mathcal{T}_i) = |\mathcal{T}_i \cap [n, n+1)|$ *for all* n*, where* $X \sim$ *Uniform*$[0,1]$*. We say that the policy* $\boldsymbol{\sigma}$ *is ordered laminar if* $\boldsymbol{\sigma}^l$ *is ordered laminar for all* $l \in [L]$*.*

For any given $l \in [L]$, an ordered laminar policy $\boldsymbol{\sigma}^l$ can be represented using a finite set of ordered laminar intervals $\{\mathcal{I}_1, \dots \mathcal{I}_{K+1}\}$, which significantly simplifies the structure of the policy. It also implicitly implies that the signals $\{S_1, \dots, S_{K+1}\}$ and their respective laminar intervals $\{\mathcal{I}_1, \dots \mathcal{I}_{K+1}\}$ follow the same ordering, where the signals are ordered in terms of the expected wait time they imply. For simplicity, we occasionally omit explicit mention of this ordering. Next, we present our main result.

Theorem 1. *Consider the throughput maximization problem under demand variation and heterogeneous customers with private patience as presented in Section 2. There exists an optimal signaling mechanism* $\boldsymbol{\Sigma}^* = (\mathcal{S}^{[K]}, \boldsymbol{\sigma}^*)$ *such that obedience is a customer equilibrium (i.e.,* $\boldsymbol{\Sigma}^*$ *is recommendation-obedient) and the policy* $\boldsymbol{\sigma}^*$ *is ordered laminar.*

Theorem 1 characterizes the structure of the optimal signaling mechanism under customer heterogeneity (with private patience information) and demand variation. We defer the discussion on the implications of Theorem 1 to Sections 4 and 5, where we examine, respectively, the effect of customer heterogeneity (with private patience information) alone, and the effect of demand variation alone, on the structure of the optimal signaling mechanism and on customer utility. In Section 4, we also provide a detailed discussion on the implications of a policy having an ordered laminar structure.

4. Customer heterogeneity with private patience information

In this section, we study the special case where there is no variation in demand to understand better the effect of customer heterogeneity with private patience information. For ease of notation, we assume that $\mathcal{L} = \{\lambda\}$, meaning $\Lambda = \lambda$ with probability 1. We drop the argument (superscript) corresponding to the arrival rate realization, i.e., the policy is $\boldsymbol{\sigma}$, and the corresponding stationary distribution is $\boldsymbol{\pi}$.

Proposition 3. *Consider the scenario without demand variation, i.e., $\Lambda = \lambda$ with probability 1. Then, the Throughput C-MDP (4) is equivalent to the following,*

$$\max_{\boldsymbol{\sigma}} \mathbb{E}_{\boldsymbol{\sigma}}\Big[\sum_{i=1}^{K} \lambda P_i \xi(S_i)\Big] \quad \textit{s.t.} \quad \mathbb{E}_{\boldsymbol{\sigma}}\big[\xi(S_i) h_i(W)\big] \geq 0, \; \forall i \in [K].$$

Proposition 3 provides a simplification of Throughput C-MDP (4) when there is no demand variation. Notably, the second IC constraint in the Throughput C-MDP (4), namely $\mathbb{E}_{\boldsymbol{\sigma}}\big[h_i(W)\xi(S_{i+1})\big] \leq 0$ for all $i \in [K]$, is trivially satisfied under no demand variation. Note that the constraint $\mathbb{E}_{\boldsymbol{\sigma}}\big[h_i(W)\xi(S_{i+1})\big] \leq 0$ implies that the expected wait time under signal S_{i+1} must be at least c_i/b_i. When the server's objective is to maximize throughput, it tends to allow as many customers as possible to enter the system, thereby increasing the overall wait time. Consequently, the expected wait time associated with signal S_{i+1} typically exceeds the lower bound c_i/b_i. In particular, we can show that if $\mathbb{E}_{\boldsymbol{\sigma}}\big[h_i(W)\xi(S_{i+1})\big] = 0$, the server can increase the throughput by providing a signal associated with a higher arrival rate (such as S_j for some $j < i+1$) in place of signal S_{i+1} for certain states where S_{i+1} was previously provided. The proof of Proposition 3 is detailed in Appendix D.1 and relies on the LP formulation of Throughput C-MDP (4) presented in Corollary 2 in Section 6.1.

Remark 1 (Positive customer utility)**.** *Under any feasible policy $\boldsymbol{\sigma}$, a customer of type i enters the system under all signals S_j with $j \leq i$. As the expected wait time under signal S_j is at most c_j/b_j, the utility received by a type-i customer is at least $c_i - b_i \times \frac{c_j}{b_j}$, when signal S_j is provided. This implies that by keeping the customer type information private, some customer types (besides type 1) can receive strictly positive utility on average. Thus, there is an incentive for customers to keep their patience information private.*

4.1. Structure of the policy

In this section, we provide details on the structural properties of a policy with laminar structure. As mentioned before, any policy can by mapped to a partition on the real line using the representation in Eq. (5), however, the representation in Eq. (5) is be highly complex as the partitions are represented as a disjoint union of countably many subsets. The policy $\boldsymbol{\sigma}$ can be simplified if the corresponding partition $\{\mathcal{T}_i\}_{i\in[K+1]}$ has structural simplicity. An example of such a simple structure is the monotone structure.

Definition 7 (Monotone policy)**.** *A partition $\{\mathcal{T}_1, \ldots, \mathcal{T}_{K+1}\}$ of $\mathbb{R}_+$ is called a monotone if there exists $0 = x_0 \leq x_1 \leq x_2 \leq \cdots \leq x_K \leq x_{K+1} = \infty$, such that $\mathcal{T}_i = [x_{i-1}, x_i)$ for all $i \in [K+1]$, where $x_{i-1} = x_i$ implies* *that $\mathcal{T}_i = \emptyset$. A policy $\boldsymbol{\sigma}$ is called monotone (or has a monotone structure) if there exists a monotone partition $\{\mathcal{T}_1, \ldots, \mathcal{T}_{K+1}\}$ such that $\sigma(S_i|n) = \mathbb{P}(n + X \in \mathcal{T}_i) = |\mathcal{T}_i \cap [n, n+1)|$ for all n, where $X \sim$ Uniform$[0,1]$.*

An example of the monotone policy is provided in Figure 3.

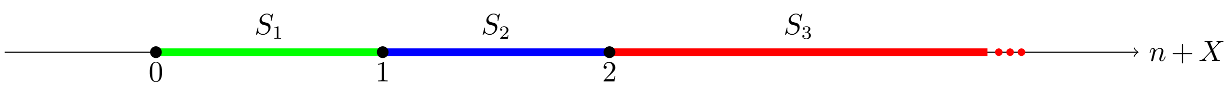

Figure 3 Pictorial representation of the example of the monotone structure with $\mathcal{T}_1 = [0,1), \mathcal{T}_2 = [1,2)$ and $\mathcal{T}_3 = [2,\infty)$

From a practical standpoint, a policy with monotone structure is desirable as it facilitates easier implementation and it can be fully characterized by a finite set of endpoints $\{x_0, x_1, \ldots, x_{K+1}\}$ to represent the policy, which significantly reduces the complexity of the policy. Note that, for the *full information* signaling mechanism, where the server discloses the current queue length, there exists an equivalent recommendation-obedient signaling mechanism (under Definition 2) with a monotone policy.

A simple example of a monotone policy is a threshold policy. It has been shown in Lingenbrink and Iyer (2019) that when customers are homogeneous ($K = 1$), a binary signaling mechanism with a threshold policy is optimal. Lingenbrink and Iyer (2019) also considers the case of heterogeneous customers with *publicly* known types, that is, the incoming customer's patience information is known to the server. In this case, a binary signaling mechanism with a type-dependent threshold policy is optimal.

A monotone structure is an extension of a threshold structure that incorporates multiple thresholds. Consider a system with two customer types ($K = 2$), where the optimal signaling mechanism consists of three signals: S_1, S_2, and S_3. A monotone structure in this context would imply that there exist thresholds x_1 and x_2 such that $\mathcal{T}_1 = [0, x_1)$, $\mathcal{T}_2 = [x_1, x_2)$ and $\mathcal{T}_3 = [x_2, \infty)$ and $\sigma(S_i|n) = \mathbb{P}(n + X \in \mathcal{T}_i)$. Such a monotone structure seems intuitive (given the result in Lingenbrink and Iyer 2019) and is commonly seen in existing literature (Chapter 4.7 of Puterman 2014, and Roy et al. 2021). However, for heterogeneous customers with private patience information, we observe that the optimal policy need not be monotone.

Example of non-monotone optimal policy: Figure 4 presents an example with two customer types, where the optimal policy is non-monotone. The server sends signal S_1 if $n + X$ lies in the green region, S_2 if in the blue region, and S_3 if in the red region. Note that the green region is nested within the blue region.

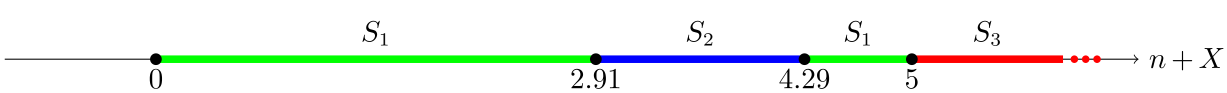

Figure 4 Pictorial representation of the optimal policy for $K = 2$ case with $p_1 = p_2 = 0.5$. Here, $c_1/b_1 = 1$ and $c_2/b_2 = 3$, and the overall arrival rate is $\lambda = 0.8$. The optimal policy is a (non-monotone) laminar policy with respect to $\{\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3\}$, where $\mathcal{T}_1 = [0, 2.91) \cup [4.29, 5)$, $\mathcal{T}_2 = [2.91, 4.29)$ and $\mathcal{T}_3 = [5, \infty)$.

When the customers' patience information is private, and their utility function is linear in W, the optimal signaling mechanism has a *laminar structure*. A (non-monotone) laminar policy is more complicated than a monotone policy, however, it is still "simple" as we only need to characterize the two endpoints for each interval $\mathcal{I}_i$ for $i \in [K+1]$. Note that a monotone partition is also a laminar partition (with $\mathcal{I}_i \cap \mathcal{I}_j = \emptyset$ for all i, j). Consequently, any monotone policy is inherently laminar. The following lemma provides a characterization of the class of ordered laminar policies.

Lemma 1. *Consider a policy $\boldsymbol{\sigma} : \mathbb{Z}_+ \to \Omega(\mathcal{S}^{[K]})$. Then, $\boldsymbol{\sigma}$ is ordered laminar if for all $i \in [K+1]$, the following holds: For any non-negative integers $\underline{n}_i$ and $\bar{n}_i$ ($\underline{n}_i < \bar{n}_i$), if $\sigma(S_i|n) > 0$ for $n = \underline{n}_i$ and $n = \bar{n}_i$, it holds that $\sigma(S_j|n) = 0$ for all $j < i$ and $\underline{n}_i < n < \bar{n}_i$.*

We illustrate Lemma 1 using the example in the right plot in Figure 2. Note that $\sigma(S_2|n) > 0$ at $n = 2$ and $n = 4$. So, according to Lemma 1, for the policy to have a laminar structure, we need $\sigma(S_1|n) = 0$ at $n = 3$, which indeed holds. The characterization presented in Lemma 1 is used in the proof of Theorem 1. In particular, we show that the optimal policy (i.e., the optimal solution of Throughput C-MDP (4)) satisfies the condition in Lemma 1.

Ordered laminar structure: Note that Lemma 1 implies a strict ordering for the nested structure of signals. In particular, for any pair of patience cutoff signals S_i and S_j with $j < i$, the queue length interval associated with S_i (which has a lower arrival rate as $P_i < P_j$) may be nested within the interval associated with S_j, but not vice versa. For example, in the optimal policy shown in Figure 4, the laminar interval corresponding to S_2 is nested within the intervals where S_1 is given.

Corollary 1. *Under heterogeneous customers with private patience information and no demand variation, there exists an optimal recommendation-obedient signaling mechanism $\boldsymbol{\Sigma}^* = (\mathcal{S}^{[K]}, \boldsymbol{\sigma}^*)$ such that $\boldsymbol{\sigma}^*$ is ordered laminar.*

Corollary 1 follows from Theorem 1, and says that there exists a laminar structure to map the queue length to the $K+1$ signals optimally. In the example presented in Figure 4, the optimal policy is laminar with $\mathcal{I}_3 = [5, \infty)$, $\mathcal{I}_2 = [2.91, 4.29)$ and $\mathcal{I}_1 = [0, 5)$. Here, $\mathcal{I}_2 \subseteq \mathcal{I}_1$. The server gives the signal "both types enter" (Signal S_1) when $n + X \in \mathcal{I}_1 \backslash \mathcal{I}_2$, "only type-2 enter" signal (Signal S_2) when $n + X \in \mathcal{I}_2$, and a "no one enters" signal (Signal S_3) when $n + X \in \mathcal{I}_3$. An interesting point to note is that the entry region for type-1 customers is divided into two parts. Suppose a type-1 customer enters the system when the state is in the "good" region, i.e., when $n + X \in [0, 2.91)$. In that case, their expected wait time is below the patience $c_1/b_1 = 1$. Conversely if they enter when the state is in the "bad" region, i.e., when $n + X \in [4.29, 5)$, their expected wait time is significantly higher than the patience. However, the overall expected wait time is equal to $c_1/b_1 = 1$, as the likelihood of the state being in the bad region is relatively low. Also, the customers do not know whether the state is in a "good" or "bad" region, which allows such a non-monotone signaling mechanism to be compatible with the customer's incentives.

Intuition on laminar structure of the optimal policy: We now illustrate the intuition behind the laminar structure of the optimal policy using the example presented in Figure 4. Suppose the server adopts a monotone policy, by deciding to provide S_1 (or "low congestion" signal) up to the largest possible queue length threshold (i.e., the expected wait time under signal S_1 is equal to $c_1/b_1 = 1$) before switching to S_2 (or "medium congestion" signal). Note that this is a sort of 'greedy' strategy. In this scenario, the server provides S_1 up to the threshold of 3.35 and S_2 in the interval $[3.35, 4)$. See Figure 5 for a pictorial illustration.

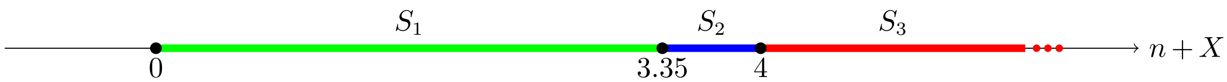

Figure 5 Pictorial representation of the optimal monotone policy under the same system parameters as in Figure 4.

If we are not restricted to the 'greedy' monotone policy, it can start issuing the blue signal at a lower threshold value than 3.35. In the optimal scenario, the server switches to the blue signal at the threshold value 2.91 and continues to provide it in the interval $[2.91, 4.29)$. Here, 4.29 is the maximum threshold up to which the server can give the blue signal while keeping the expected wait time under the blue signal less than $c_2/b_2 = 3$. Since the server stops providing the green signal at a lower threshold value than in the monotone case, the expected wait time under the green signal for the interval $[0, 2.91)$ is below $c_1/b_1 = 1$. This gives the server some slack to provide the green signal again over another small interval above 4.29, specifically in the interval $[4.29, 5)$. By doing so, the server can allow customer entry over a larger interval, i.e., up to 5 compared to 4 in the monotone case, which might increase the throughput. Overall, using a policy with a non-monotonic structure might increase throughput by allowing customer entry over a larger interval for $W + X$ while still satisfying the incentive compatibility constraints.

4.2. Empirical comparison of signaling mechanisms under customer heterogeneity

We consider four kinds of signaling mechanisms:

(i) *No information:* The server provides no information or signal to any incoming customer. This signaling mechanism results in a customer equilibrium where a customer enters the system with a type-dependent probability, say ν_i, for $i \in [K]$. One can easily compute the value of ν_i's, and conclude that, if $\nu_i > 0$, then $\nu_j = 1$ for all $j > i$.

(ii) *Full information:* The server reveals the exact current queue length to each incoming customer. Under full information, the customer of type i enters if the queue length is less than or equal to c_i/b_i, and leaves otherwise. Note that there exists a recommendation-obedient signaling mechanism with a monotone policy that is equivalent (under Definition 2) to the full information signaling mechanism.

(iii) *Monotone signaling mechanisms:* We consider the class of recommendation-obedient signaling mechanisms with a monotone policy. Note that there does not exist an LP formulation to solve for optimal monotone policy. We use a gradient descent algorithm to search for the optimal monotone policy.

(iii) *Laminar signaling mechanisms:* We consider the class of recommendation-obedient signaling mechanisms with a general laminar policy. From Corollary 1, we know this will give us an optimal policy. Further, the optimal policy can be computed using the LP formulation of Throughput C-MDP (4) presented in Section 6.1.

In Figure 6, we consider two examples with varying levels of arrival rate, and for both examples, we provide the throughput for the four kinds of signaling mechanisms mentioned above. We also provide the structure of the optimal policy and the optimal monotone policy for the two examples in Figure 7.

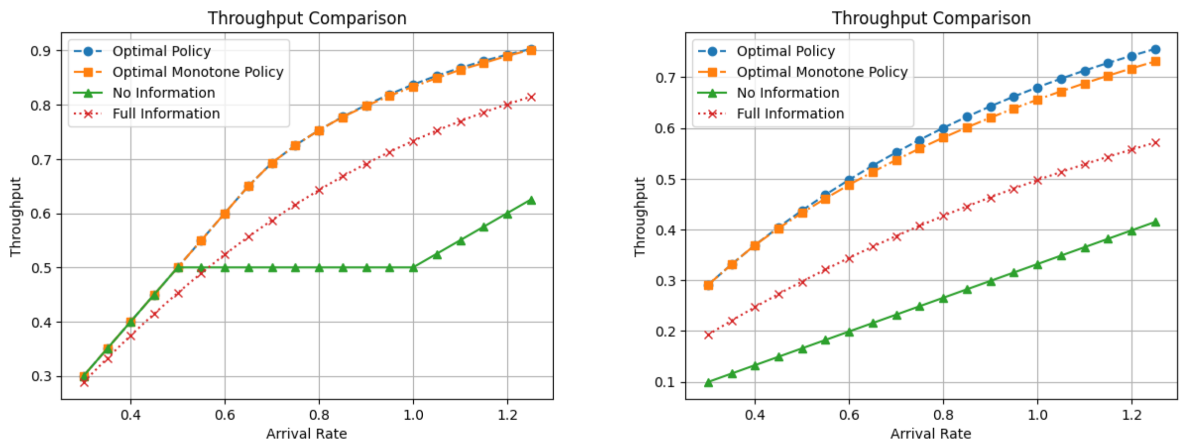

Figure 6 The comparison of throughput with respect to the arrival rate for different signaling mechanisms. In the first plot, $K = 2$ with $p_1 = p_2 = 0.5$, and $c_1 = 2$ and $c_2 = 4$. In second plot, $K = 3$ with $p_1 = 0.38, p_2 = 0.28$ and $p = 0.33$, and $c_1 = 0.29$ and $c_2 = 0.97$ and $c_3 = 2.37$. In both cases $b_i = 1$ for all i.

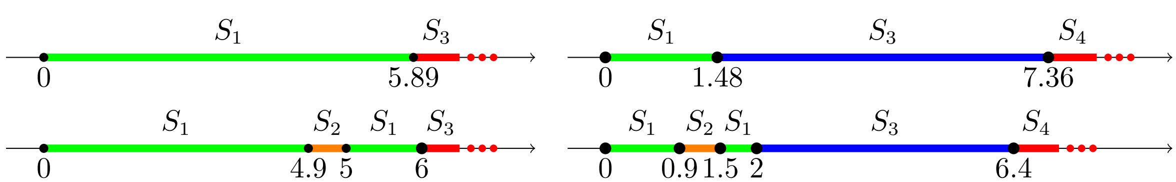

Figure 7 The structure of the optimal monotone policy (above) and the optimal policy (below) for the two examples considered in Figure 6 with the overall arrival rate $\lambda = 0.85$. (The plot is not to scale.)

In both plots presented in Figure 6, we observe that the optimal mechanism significantly outperforms both the no-information and full-information mechanisms. In the left plot, we also notice that the throughput achieved under the optimal monotone policy is quite close to that of the optimal policy even though the optimal policy exhibits a non-monotonic laminar structure for some arrival rate values (e.g., $\lambda = 0.85$ as seen in Figure 7). In contrast, the right plot in Figure 6 presents a case with a noticeable difference between the throughput of the optimal monotone policy and that of the optimal policy. Furthermore, as illustrated in Figure 7, the optimal policy can look quite different from the optimal monotone policy; notably, in the monotone policy, the server chooses not to provide the signal S_2 at all. Also, note that in both the plots in Figure 7, the queue length interval where the server provides signal S_1 or S_2 is larger under the optimal policy than in the monotone policy case.

5. Demand variation

This section looks at the second feature considered in this work, variable demand. As explained in Section 2, the arrival rate varies from period to period and follows the distribution $\mathcal{A}$, i.e., $\Lambda \sim \mathcal{A}$. With a little abuse of notation, we use λ to denote the realized arrival rate. We assume λ is the server's private information, meaning customers do not know the realized arrival rate λ (they only know the distribution $\mathcal{A}$). In this case, the system state is given by $(W, \Lambda) = (n, \lambda)$, which includes the queue length and the arrival rate realization.

To gain insights into the effect of variable demand on the signaling mechanism, we consider the case when the customers are homogeneous ($K = 1$), i.e., the utility function of each customer is given by $h(n) = c - bn$, where c and b are constants. In this case, our results in Section 3.3 imply that for a fixed arrival rate, there exists an optimal signaling mechanism where the server uses a binary signal ("Enter" (1) or "Leave" (0)) with a threshold. For random arrival rate, the optimal signaling mechanism is still binary signaling with a λ-dependent threshold structure.

Proposition 4 (Binary signaling with threshold policy)**.** *Under homogeneous customers and demand variation, there exists an optimal signaling mechanism* $\mathbf{\Sigma} = (\mathcal{S}^{[1]}, \boldsymbol{\sigma})$ *with* binary *patience cutoff signals (i.e.,* $\mathcal{S}^{[1]} = \{1, 0\}$*), such that obedience is a customer equilibrium and the optimal policy* $\boldsymbol{\sigma}$ *is a threshold policy, i.e., for each* $l \in [L]$*, there exists* T_l *such that* $\sigma^l(S_1|n) = \mathbb{P}\big(n + X \in [0, T_l)\big)$*, where* $X \sim$ *Uniform*$[0, 1]$.

Proposition 4 states that when the customers are homogeneous, if the realized arrival rate is λ_l, the server recommends that the customer enter the queue if $n + X < T_l$, and recommends that the customer leave otherwise. The set of thresholds $\{T_l\}_{l \in [L]}$ is chosen based on the distribution $\mathcal{A}$ and the utility function $h(\cdot)$. The result in Proposition 4 follows by using the result in Theorem 1 (under $K = 1$), and showing that if the optimal laminar policy is non-monotone, it can be perturbed to create a threshold policy which satisfies the IC constraints and achieves the same throughput. The proof is provided in Appendix D.3.

5.1. Private vs public knowledge of realized arrival rate

Let $\text{Th}^*(\lambda)$ denote the throughput under the optimal signaling mechanism when the arrival rate is fixed and given by λ, i.e., when $\Lambda = \lambda$ with probability 1. Similarly, $\text{Th}^*(\mathcal{A})$ denote the optimal throughput when $\Lambda \sim \mathcal{A}$. In this section, we compare two scenarios:

(i) *Scenario 1 [The realized arrival rate is public knowledge]:* In this scenario, we assume that both the server and the customers know the arrival rate realization. This scenario is motivated by the fact that, in some cases, the customers are aware of and can observe the features that affect the demand fluctuations. For example, the congestion level in many service systems (e.g., online restaurant reservations) may vary depending on whether it is a weekday or a weekend. Both the server and the customers can be aware of these fluctuations. Another example is when the server vows to be transparent with customers and provides the arrival rate information as part of the signal.

In this scenario, if the realized arrival rate is $\Lambda = \lambda$, the server uses the optimal signaling mechanism corresponding to the system for which the arrival rate is fixed and given by λ. As such, the overall throughput is given by $\mathbb{E}_{\Lambda\sim\mathcal{A}}[\text{Th}^*(\Lambda)]$.

(ii) *Scenario 2 [The realized arrival is the server's private information]:* This is the scenario presented in Section 2 where only the server knows the arrival rate realization. In this scenario, the maximum achievable throughput is $\text{Th}^*(\mathcal{A})$.

Proposition 5. *For homogeneous customers with a linear utility function, i.e., $h(n) = c - bn$, $\mathbb{E}_{\Lambda\sim\mathcal{A}}[Th^*(\Lambda)] \leq Th^*(\mathcal{A})$.*

Proposition 5 states that if the customer is unaware of the arrival rate realization, the server benefits from keeping the information private, i.e., the server can exploit the information this asymmetry to (possibly) increase the system throughput. Section 5.2 provides an empirical comparison showcasing a case where the inequality is strict. The proof of the proposition is straightforward and uses the fact that the server can always reveal the arrival rate to the customer, implying that if $\mathbb{E}_{\Lambda\sim\mathcal{A}}[\text{Th}^*(\Lambda)]$ is achievable system throughput in scenario 1, then it is also achievable in scenario 2. Mathematically, we compare the feasibility sets of the LPs in the two scenarios. For completeness, the proof of Proposition 5 is provided in Appendix D.2.

Flexibility in the thresholds. The customer's lack of arrival rate information provides the server with extra flexibility, compared to the case when customers know the arrival rate realization. In particular, the server can increase the expected wait times on days with some arrival rates and compensate by decreasing them for other arrival rates to boost overall throughput. To illustrate this, consider a simple scenario where Λ take two values, $\Lambda \sim \text{Unif}(\{\lambda_L, \lambda_H\})$, i.e., $\Lambda = \lambda_L$ with probability 0.5 and $\Lambda = \lambda_H$ with probability 0.5, where $\lambda_L < \lambda_H$. When the arrival rate realization is public knowledge, the server uses thresholds T_L and T_H for $\Lambda = \lambda_L$ and $\Lambda = \lambda_H$, respectively. On the other hand, when the arrival rate realization is only known to the server, different thresholds $\hat{T}_L$ and $\hat{T}_H$ are used.

The server sets these thresholds such that on a low demand day ($\Lambda = \lambda_L$), it allows more customers to enter than it would if the customers knew that $\Lambda = \lambda_L$, i.e., $\hat{T}_L > T_L$. However, this results in longer wait times. To compensate, on high demand days ($\Lambda = \lambda_H$), the server lowers the entry threshold ($\hat{T}_H < T_H$). By carefully choosing $\hat{T}_L$ and $\hat{T}_H$, the server can ensure that the overall expected wait time remains unchanged while increasing the overall throughput.

Remark 2. *From the above discussion, we can also conclude that when the arrival rate realization is the server's private information (as opposed to being public knowledge), the expected wait time experienced by the customers on low demand days might be higher than the expected wait time on high demand days. Conversely, the probability of entry is higher on low-demand days and lower on high-demand days. This phenomenon showcases the effect of variable demand when there is a misalignment between the server's objective and the customer's utility.*

5.2. Empirical comparison

In this section, we provide an empirical comparison of the system throughput as a function of the variability in demand. We consider a simple example where $\Lambda \sim \text{Unif}(\{\lambda_L, \lambda_H\})$. Define $\bar{\lambda} = \frac{1}{2}(\lambda_L + \lambda_H)$ and $\delta = \frac{1}{2}(\lambda_H - \lambda_L)$. In Figure 8, we use $\bar{\lambda} = 2$. In addition to the two scenarios considered in Section 5.1: (i) "δ-public," the realized arrival rate is public knowledge, and (ii) "δ-private," the realized arrival rate is the server's private information, we also consider a third scenario: "$\delta = 0$", where the arrival rate is fixed and given by $\bar{\lambda}$. We use the $\delta = 0$ case as a benchmark to demonstrate the effect of variation in demand on system throughput and the corresponding thresholds.

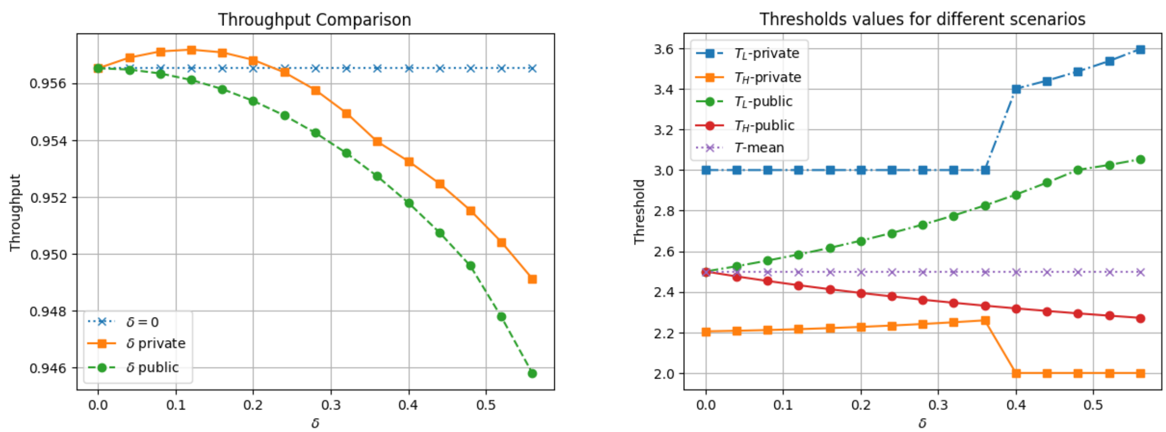

Figure 8 The effect of demand variation (i.e., with respect to δ) on the throughput and corresponding thresholds for the three considered scenarios. Here, $h(n) = 2 - n$, $\lambda \sim \text{Unif}(\{\lambda_L, \lambda_H\})$ and $\bar{\lambda} = 2$.

Observations on throughput: The first plot in Figure 8 compares the overall throughput for the three scenarios. In line with the result in Proposition 5, the δ-private scenario achieves higher throughput than the δ-public scenario. We note that while the inequality in Proposition 5 is weak, the throughput gains from the arrival rate being privately known to the server are strictly positive in this example.

We also note that, compared to the deterministic arrival rate (i.e., $\delta = 0$) case, demand variation can lead to a higher throughput in certain situations, i.e., there are examples where $\text{Th}^*(\bar{\lambda}) \le \text{Th}^*(\mathcal{A})$. In Figure 8, the server attains a higher throughput in the δ-private scenario compared to the $\delta = 0$ scenario for $\delta \in (0, 0.22)$. Overall, our simulation results indicate that a small amount of demand fluctuation provides the server with more flexibility, facilitating throughput gains. However, large fluctuations result in a loss of overall throughput due to significant under-utilization when the arrival rate is low. Note that while small variability in demand can improve the system throughput, the server cannot control the variability itself.

Observations on thresholds: In the second plot of Figure 8, we observe how the thresholds change in the δ-public and δ-private scenarios in response to variations in demand, δ. In the δ-public scenario, T_L-public denote the threshold on low-demand days ($\Lambda = \lambda_L$) and T_H-public denote the threshold on high-demand days ($\Lambda = \lambda_H$). Similarly, we define T_L-private and T_H-private for the δ-private case. Additionally, let T-mean denote the threshold when $\delta = 0$. Note that, in line with our discussion on flexibility in the thresholds in Section 5.1 and Remark 2, we see that T_L-public $> T_L$-private and T_H-public $< T_H$-private for all $\delta > 0$.

Another observation here is that in the δ-private case, the server employs two very distinct thresholds (T_L-private and T_H-private) even when $\delta = 0$. The reason behind this phenomenon is as follows: a real-valued threshold is analogous to a randomized policy, and there are multiple ways to implement such a policy. Suppose that $\delta = 0$. In the δ-private case depicted in the second plot of Figure 8, the server can implement the randomized policy corresponding to the $\delta = 0$ case by alternating between a large threshold on some days (e.g., T_L-private on odd-numbered days) and a smaller threshold on other days (e.g., T_H-private on even-numbered days). Mathematically, when solving for the optimal solution in the δ-private case using a linear program with $\delta = 0$, there are multiple solutions for T_L-private and T_H-private. When $\delta > 0$ but small, the optimal thresholds in the δ-private case are unique, and the thresholds T_L-private and T_H-private differ significantly from each other. This allows the server to achieve higher throughput in the δ-private case compared to the $\delta = 0$ case for small but positive values of δ.

6. Mathematical details and proof of Theorem 1

In this section, we provide the mathematical machinery underlying our methodology and the key arguments behind the proof of Theorem 1. In Section 6.1, we present the LP formulation for the C-MDP (4). Section 6.2 introduces the Lagrangian dual of the C-MDP (4), which we in turn use to prove Theorem 1, as detailed in Section 6.3. Before delving into the details, we provide the intuition behind the proof of Theorem 1.

Intuition behind the proof of Theorem 1: To prove Theorem 1, we establish a Lagrangian dual of the C-MDP (4). This requires proving that strong duality holds, which includes two steps. First, we map the set of feasible policies to a compact and convex set of variables, such that the objective function and the IC constraints are affine functions of the variables. Next, we identify a policy that satisfies Slater's condition. The combination of the two steps shows that strong duality holds. More details are provided in Section 6.2.

Using strong duality, we consider the Lagrangian relaxation of the problem (4) under the optimal Lagrange multipliers, which is simply a continuous-time average cost problem. By analyzing the corresponding Bellman optimality equation, we establish the convexity of the value function (see Lemma 3). Finally, using the convexity result, we show that any policy satisfying the Bellman optimality equation and the complimentary slackness conditions is laminar, which implies that the policy under the optimal signaling mechanism is laminar.

6.1. Linear programming formulation

The C-MDP (4) can be written as a linear program in terms of the mapping $\boldsymbol{\sigma}$ and the corresponding tuple of steady-state distributions $\boldsymbol{\Pi}$. To do that, we define the variable $\boldsymbol{\phi}$ as follows.

$$\phi_{n,i}^l := \mathbb{P}_{\boldsymbol{\sigma}}(W = n, \xi(S_i) = 1 | \Lambda = \lambda_l) = \pi_n^l \sigma^l(S_i|n).$$

Further, for all $l \in [L]$, $\boldsymbol{\phi}_i^l = \{\phi_{n,i}^l\}_{n\geq 0}$ and $\boldsymbol{\phi}^l = \{\boldsymbol{\phi}_1^l, \boldsymbol{\phi}_2^l, \ldots, \boldsymbol{\phi}_{K+1}^l\}$. Finally, $\boldsymbol{\phi} = \{\boldsymbol{\phi}^1, \boldsymbol{\phi}^2, \ldots, \boldsymbol{\phi}^L\}_{l\in[L]}$. Note that one can recover $\boldsymbol{\Pi}$ and $\boldsymbol{\sigma}$ from the variable $\boldsymbol{\phi}$ using the following:

$$\pi_n^l = \sum_{i=1}^{K+1} \phi_{n,i}^l, \text{ and } \sigma^l(S_i|n) = \frac{\phi_{n,i}^l}{\pi_n^l} \text{ if } \pi_n^l > 0,$$

and if $\pi_n^l = 0$, set $\sigma^l(S_{K+1}|n) = 1$ and $\sigma^l(S_i|n) = 0$ for $i \in [K]$.

Corollary 2. *The throughput maximization problem* (4) *is equivalent to the following infinite dimensional linear program.*

$$\begin{aligned}
\textit{Throughput LP}: \quad \max_{\boldsymbol{\phi}} \quad & \sum_{i=1}^{K}\sum_{l=1}^{L} q_l \lambda_l P_i \langle \boldsymbol{\phi}_i^l, \mathbb{1}\rangle \\
\textit{s.t.} \quad & \sum_{l=1}^{L} q_l \langle \boldsymbol{\phi}_i^l, \mathbf{h}_i\rangle \geq 0, \ \forall i \in [K], \\
& \sum_{l=1}^{L} q_l \langle \boldsymbol{\phi}_{i+1}^l, \mathbf{h}_i\rangle \leq 0, \ \forall i \in [K], \\
& \lambda_l \sum_{i=1}^{K} P_i \phi_{n,i}^l - \sum_{i=1}^{K+1} \phi_{n+1,i}^l = 0, \ \forall n \geq 0, \textit{ and } \forall l \in [L], \\
& \sum_{i=1}^{K+1} \langle \boldsymbol{\phi}_i^l, \mathbb{1}\rangle = 1, \ \forall l \in [L], \\
& \phi_{n,i}^l \geq 0, \ \forall n \geq 0 \textit{ and } \forall i \in [K], \forall l \in [L].
\end{aligned} \tag{6}$$

where $P_i = \sum_{j=i}^{K} p_j$ *and* $P_{K+1} = 0$, *and* $\mathbf{h}_i := \{h_i(n)\}_{n\geq 0} = \{c_i - b_i n\}_{n\geq 0}$. *Suppose* $\boldsymbol{\phi}^*$ *is an optimal solution of the LP in* (6)*. Then, the optimal signaling scheme is given by* $\sigma^{*,l}(S_i|n) = \frac{\phi_{n,i}^{*,l}}{\sum_{j=1}^{K+1} \phi_{n,j}^{*,l}}$.

Corollary 2 reduces the seemingly complex optimization problem (4) to a much simpler linear program. The LP formulation (6) in Corollary 2 follows simply by substituting the variable $\boldsymbol{\phi}$ in the constrained MDP formulation in (4). In the objective, for a given $l \in [L]$, the term $\langle \boldsymbol{\phi}_i^l, \mathbb{1}\rangle$ is the total probability of giving signal S_i, and $\lambda_l P_i$ is the cumulative arrival rate under signal S_i. As such, $\sum_{i=1}^{K}\sum_{l=1}^{L} q_l \lambda_l P_i \langle \boldsymbol{\phi}_i^l, \mathbb{1}\rangle$ is the throughput of the system. The constraint $\sum_{l=1}^{L} q_l \langle \boldsymbol{\phi}_i^l, \mathbf{h}_i\rangle \geq 0$, $\forall i \in [K]$ comes from the first IC constraint in (4), capturing that the expected wait time given signal S_i is at most c_i/b_i. Similarly, the second IC constraint in (4) gives $\sum_{l=1}^{L} q_l \langle \boldsymbol{\phi}_{i+1}^l, \mathbf{h}_i\rangle \leq 0$, $\forall i \in [K]$, denoting that the expected wait time given signal S_i is at least c_{i-1}/b_{i-1}. The next set of equality constraints captures the detailed balance conditions of the Markov chain for each l. The last set of equality constraints captures that the steady state distribution of the Markov chain must sum to 1, for each l. Finally, since the $\phi_{n,i}^l$'s are probabilities, they must be non-negative.

Note that the Throughput LP (6) is an infinite dimensional linear program. One can obtain an approximate solution by truncating the number of variables in the linear program, which corresponds to considering the system with a finite buffer. For more details on solving infinite dimensional LPs, please refer to Nash and Anderson (1987). After solving the LP (6), we can obtain the optimal policy. The LP formulation not only provides a computationally efficient way to solve for the optimal policy, but it is also used in the proof of Proposition 3 and Theorem 1.

6.2. Lagrangian approach

In this section, we refer to signals as actions, and $\boldsymbol{\sigma} = \{\boldsymbol{\sigma}^l\}_{l\in[L]}$ denotes the policy. For the C-MDP (4), we denote the primal problem and corresponding dual problem as

$$\text{Primal:} \quad \max_{\boldsymbol{\sigma}} \min_{\tau_i,\tilde{\tau}_i\geq 0} H(\boldsymbol{\sigma};\boldsymbol{\tau},\tilde{\boldsymbol{\tau}}) \qquad\qquad \text{Dual:} \quad \min_{\tau_i,\tilde{\tau}_i\geq 0} \max_{\boldsymbol{\sigma}} H(\boldsymbol{\sigma};\boldsymbol{\tau},\tilde{\boldsymbol{\tau}}), \tag{7}$$

where we define the Lagrangian function $H(\boldsymbol{\sigma};\boldsymbol{\tau},\tilde{\boldsymbol{\tau}}) := \mathbb{E}_{\boldsymbol{\sigma}}\left[\sum_{i=1}^{K+1}\xi(S_i)H_i(W,\Lambda;\tau_i,\tilde{\tau}_i)\right]$, and

$$H_i(n,\lambda;\tau_i,\tilde{\tau}_i) = \begin{cases} \lambda P_1 + \tau_1 h_1(n) & \text{for } i=1 \\ \lambda P_i + \tau_i h_i(n) - \tilde{\tau}_i h_{i-1}(n) & \text{for } 1<i<K+1 \\ -\tilde{\tau}_{K+1}h_K(n) & \text{for } i=K+1. \end{cases}$$

Let $\boldsymbol{\sigma}^*$ denote an (arbitrary) optimal solution of the primal problem, i.e., it is an optimal policy for (4), and $(\boldsymbol{\tau}^*,\tilde{\boldsymbol{\tau}}^*)$ denote an optimal solution of the dual, i.e., it is the optimal solution of $\min_{\tau_i,\tilde{\tau}_i\geq 0} g(\boldsymbol{\tau},\tilde{\boldsymbol{\tau}})$ where $g(\boldsymbol{\tau},\tilde{\boldsymbol{\tau}}) := \max_{\boldsymbol{\sigma}} H(\boldsymbol{\sigma};\boldsymbol{\tau},\tilde{\boldsymbol{\tau}})$.

Lemma 2. *The optimization problem* (4) *satisfies strong duality, i.e.,*

$$\max_{\boldsymbol{\sigma}} \min_{\tau_i,\tilde{\tau}_i\geq 0} H(\boldsymbol{\sigma};\boldsymbol{\tau},\tilde{\boldsymbol{\tau}}) = \min_{\tau_i,\tilde{\tau}_i\geq 0} \max_{\boldsymbol{\sigma}} H(\boldsymbol{\sigma};\boldsymbol{\tau},\tilde{\boldsymbol{\tau}}).$$

This implies that any optimal policy $\boldsymbol{\sigma}^*$ *and any optimal Lagrange multipliers* $(\boldsymbol{\tau}^*,\tilde{\boldsymbol{\tau}}^*)$ *satisfy* $\boldsymbol{\sigma}^* \in \arg\max_{\boldsymbol{\sigma}} H(\boldsymbol{\sigma};\boldsymbol{\tau}^*,\tilde{\boldsymbol{\tau}}^*)$ *and* $(\boldsymbol{\tau}^*,\tilde{\boldsymbol{\tau}}^*) \in \arg\min_{\tau_i,\tilde{\tau}_i\geq 0} H(\boldsymbol{\sigma}^*;\boldsymbol{\tau},\tilde{\boldsymbol{\tau}})$*, and*

$$\tau_i^*\mathbb{E}_{\boldsymbol{\sigma}^*}\left[\xi(S_i)h_i(W)\right] = \tilde{\tau}_{i+1}^*\mathbb{E}_{\boldsymbol{\sigma}^*}\left[\xi(S_{i+1})h_i(W)\right] = 0,\ \forall i\in[K], \tag{8}$$

i.e., complementary slackness holds. Also, for any $i\in[K]$*, either* $\mathbb{E}_{\boldsymbol{\sigma}^*}\left[\xi(S_i)\right]=0$ *or we have* $\tau_i^*\tilde{\tau}_i^*=0$*.*

For any $i\in[K]$ such that $\mathbb{E}_{\boldsymbol{\sigma}^*}\left[\xi(S_i)\right]>0$, the condition $\tau_i^*\tilde{\tau}_i^*=0$ holds simply because we cannot have $\mathbb{E}_{\boldsymbol{\sigma}^*}\left[\xi(S_i)h_{i-1}(W)\right] = \mathbb{E}_{\boldsymbol{\sigma}^*}\left[\xi(S_i)h_i(W)\right]=0$, and so, by complementary slackness, we get $\tau_i^*\tilde{\tau}_i^*=0$.

The idea behind the proof of Lemma 2 is as follows. We consider a set $\mathcal{X}\subseteq\mathbb{R}^{2(K+1)\times L}$ that is obtained by mapping each feasible policy $\boldsymbol{\sigma}$ to the set of variables $\left\{\mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_i)|\Lambda=\lambda_l],\mathbb{E}_{\boldsymbol{\sigma}}[W\xi(S_i)|\Lambda=\lambda_l]\right\}_{i\in[K+1],l\in[L]}$. Then, the optimization problem (4) can be mapped to a maximization problem with linear objective function and linear constraints, where the variables lie in the set $\mathcal{X}$. We show that the set $\mathcal{X}$ is closed, convex, and bounded. In the final step, we show that Slater's condition holds by constructing a policy for which the IC constraints are strictly satisfied. The complete proof of Lemma 2 is provided in Appendix E.1.

As previously discussed, we use the C-MDP formulation as an analytical tool to derive the structural properties of the optimal policy. A critical step in our analysis is the application of strong duality. Specifically, we establish structural properties of the optimal value function of problem $\max_{\boldsymbol{\sigma}} H(\boldsymbol{\sigma};\boldsymbol{\tau}^*,\tilde{\boldsymbol{\tau}}^*)$ to characterize the structure of $\boldsymbol{\sigma}^*$.

6.3. Proof of Theorem 1

In this section, we provide the proof of Theorem 1, with some of the more detailed arguments (i.e., the proof of Lemma 3 and Lemma 4) deferred to the appendix. Based on Lemma 2, we study $\max_{\boldsymbol{\sigma}} H(\boldsymbol{\sigma};\boldsymbol{\tau}^*,\tilde{\boldsymbol{\tau}}^*)$, which is a standard continuous-time average cost problem (Guo et al. 2009, Chapter 7 and Chapter 11). In particular, we show that any policy satisfying the conditions outlined in Lemma 2 has a laminar structure. For convenience, we drop the superscript in * in $\boldsymbol{\sigma}^*$ and consider $\boldsymbol{\sigma}$ to be the optimal policy. We denote b^* as the corresponding optimal value, i.e., $b^* = \max_{\boldsymbol{\sigma}} H(\boldsymbol{\sigma};\boldsymbol{\tau}^*,\tilde{\boldsymbol{\tau}}^*)$. We also define $H_i^l(n) := H_i(n,\lambda_l;\tau_i^*,\tilde{\tau}_i^*)$ for all $i \in [K+1]$ and $l \in [L]$. Without loss of generality, we assume that if $\sigma^l(S_{K+1}|n) = 1$ for some n, then $\sigma^l(S_{K+1}|n') = 1$ for all $n' \geq n$. This is because $\sigma^l(S_{K+1}|n) = 1$ implies that $\pi^l_{n'} = 0$ for all $n > n'$ and so choosing $\sigma^l(S_{K+1}|n') = 1$ for all $n' \geq n$ gives us an equivalent policy.

In the first step, we assume that, for the optimal policy $\boldsymbol{\sigma}$, there exists n_0 and $\epsilon > 0$ such that $\sum_{i=1}^K \lambda_l P_i \sigma^l(S_i|n) < 1-\epsilon$ for all $n \geq n_0$ and for all $l \in [L]$. This is a strong stability assumption, and it implies that the optimal policy satisfies Bellman optimality equation provided below (see Guo et al. 2009, for more details). For a given $l \in [L]$, let $G^l(n'|n,i)$ denote the transition rate from state n to state n' upon taking action S_i, given as

$$G^l(n'|n,i) = \lambda_l P_i \mathbb{I}_{\{n'=n+1\}} + \mathbb{I}_{\{n'=n-1,n>0\}} - \big(\lambda_l P_i + \mathbb{I}_{\{n>0\}}\big)\mathbb{I}_{\{n'=n\}},$$

and $V^l(n)$ denotes the relative cost function for the state (n,λ_l). For any function $f:\mathbb{N}\to\mathbb{R}$, we define $\Delta f(n) := f(n+1) - f(n)$. Then, for all $l \in [L]$,

$$\begin{aligned} b^* &= \max_{i\in[K+1]} \Big\{ H_i^l(n) + \sum_{n'=0}^{\infty} G^l(n'|n,i) V^l(n') \Big\} \\ &= \max_{i\in[K+1]} \Big\{ H_i^l(n) + \lambda_l P_i V^l(n+1) + V^l(n-1)\mathbb{I}_{\{n>0\}} - (\lambda_l P_i + \mathbb{I}_{\{n>0\}}) V^l(n)] \Big\} \\ &= \max_{i\in[K+1]} \Big\{ H_i^l(n) + \lambda_l P_i \Delta V^l(n) \Big\} - \Delta V^l(n-1)\mathbb{I}_{\{n>0\}}\,. \end{aligned} \tag{9}$$

Here the first equality follows by using the average cost optimality equation for continuous-time Markov decision processes, see Guo et al. (2009, Chapter 5 and Chapter 7) for more details.

Further, for any randomized optimal policy, $\sigma^l(S_i|n) > 0$ only if $i \in \arg\max_{i\in[K+1]} \Big\{ H_i^l(n) + \lambda_l P_i \Delta V^l(n) \Big\}$, as otherwise one can find a policy with higher objective value. For more details, see Sutton and Barto (2018, Chapter 4). As such, for any optimal policy, we also have,

$$b^* = \sum_{i=1}^{K+1} \sigma^l(S_i|n) \Big[H_i^l(n) + \lambda_l P_i \Delta V^l(n) \Big] - \Delta V^l(n-1)\mathbb{I}_{\{n>0\}}. \tag{10}$$

For simplicity, we define $\beta^l(n) := \Delta V^l(n)$. Thus, by using Eq. (9) and Eq. (10), for all $l \in [L]$ and $n \geq 1$,

$$\beta^l(n-1) = \max_{i\in[K+1]} \big\{ \lambda_l P_i \beta^l(n) + H_i^l(n) \big\} - b^* = \sum_{i=1}^{K+1} \sigma^l(S_i|n) \Big[\lambda_l P_i \beta^l(n) + H_i^l(n) \Big] - b^*, \tag{11}$$

and at $n=0$, we have $b^* = \max_{i\in[K+1]}\left\{\lambda_l P_i \beta^l(0) + H_i^l(0)\right\}$. Next, we show that the function $\beta^l(n)$ is weakly convex.

Lemma 3. *Denote $\Delta^2\beta^l(n) := \Delta\beta^l(n+1) - \Delta\beta^l(n)$. For all $l \in [L]$, for any $n \geq 1$ and any $i \in [K+1]$ such that $\sigma^l(S_i|n+1) > 0$, we have that*

$$\Delta^2\beta^l(n-1) \geq \lambda_l P_i \Delta^2\beta^l(n)\,. \tag{12}$$

Furthermore, $\beta^l(n)$ is weakly convex, i.e., $\Delta^2\beta^l(n) \geq 0$ for all $n \geq 0$.

The proof of Lemma 3 relies on the "max" operator and linearity of $H_i^l(\cdot)$ in Eq. (11), and is provided in Appendix E.2. Similar ideas have also been used in (Altman 2021, Chapter 5).

Note that Eq. (12) implies that if $\beta^l(\cdot)$ is strictly convex for any n, i.e., $\Delta^2\beta^l(n) > 0$, then for all $n' < n$, we have $\Delta^2\beta^l(n') > 0$. Consequently, the $\beta^l(\cdot)$ function can be divided into two parts: up to a certain threshold in n, the function is strictly convex, and above that threshold, the function is linear. Mathematically, there exists N^l (possibly infinite), such that for all $n < N^l$, $\Delta^2\beta^l(n) > 0$, and $\Delta^2\beta^l(n) = 0$ for all $n \geq N^l$. We exploit this result to prove the next lemma, which directly implies that the optimal policy is laminar based on the characterization in Lemma 1.

Lemma 4. *Suppose N_1 and N_2 are queue length values such that $\sigma^l(S_i|n) > 0$ for $n \in \{N_1, N_2\}$, then, for all $N_1 < n < N_2$, we have that $\sigma^l(S_j|n) = 0$ for all $j < i$.*

We now present a sketch of the proof of Lemma 4. For action S_i and a given $l \in [L]$, we define the action value function for S_i at queue length n to be $Q^l(n, S_i) := \lambda_l P_i \beta^l(n) + H_i^l(n) - b^*$. As such, from Eq. (11), $\beta^l(n-1) = \max_{i\in[K+1]} Q^l(n, S_i)$. And for any pair of signals (S_i, S_j), by definition of $Q^l(n,\cdot)$ we have

$$Q^l(n, S_i) - Q^l(n, S_j) = \lambda_l (P_i - P_j)\beta^l(n) + H_i^l(n) - H_j^l(n)\,. \tag{13}$$

Note that, from Eq. (11), at any $n \geq 0$, $\sigma^l(S_i|n) > 0$ only if $Q^l(n, S_i) = \max_j\{Q^l(n, S_j)\}$. This implies that for any choice of i and j with $j < i$, if $\sigma^l(S_i|n) > 0$ for $n \in \{N_1, N_2\}$, then $Q^l(n, S_i) \geq Q^l(n, S_j)$ for $n \in \{N_1, N_2\}$. Further, for any $j < i$, the difference $Q^l(n, S_i) - Q^l(n, S_j)$ is weakly concave in terms of n, as $\Delta^2[Q^l(n, S_i) - Q^l(n, S_j)] = \lambda_l (P_i - P_j)\Delta^2\beta^l(n) \leq 0$ by Lemma 3 and $P_i < P_j$.

If $\beta^l(n)$ is *strictly* concave for $n = N_1$, i.e., $\Delta^2\beta^l(n) > 0$ for $n = N_1$, by concavity of $Q^l(n, S_i) - Q^l(n, S_j)$ for $n = N_1$ (and some additional technical arguments), we get that for all $N_1 < n < N_2$, $Q^l(n, S_i) > Q^l(n, S_j)$, which implies that S_i strictly dominates S_j for all such n. On the other hand, if $\beta^l(n)$ is linear (i.e., $\Delta^2\beta^l(n) = 0$) for $n = N_1$, then, we show that there is a unique action S_i such that $Q^l(n, S_i) = \max_j\{Q^l(n, S_j)\}$, i.e., there cannot be multiple actions with same value function. Now, by using Lemma 1, we get that $\boldsymbol{\sigma}^l$ is ordered laminar for all $l \in [L]$.

In the last step, we verify that our initial assumption, that there exists n_0 and ϵ such that $\sum_{i=1}^{K} \lambda_l P_i \sigma^l(S_i|n) < 1 - \epsilon$ for all $n \geq n_0$ and for all $l \in [L]$. For any $l \in [L]$ and $i \in [K+1]$, suppose N_i^l is the smallest (possibly infinite) queue length value such that $\sigma^l(S_i|n) = 0$ for all $n > N_i^l$. As $\boldsymbol{\sigma}^l$ is laminar for all $l \in [L]$, we also conclude that, for any $l \in [L]$, there is a unique signal S_{i_l}, such that $N_{i_l}^l = \infty$ and for any $j \neq i_l$, $N_j^l < \infty$, as otherwise, the statement in Lemma 1 will be violated. Let $n_0^l = \max_{j \neq i_l} N_j^l$. Then, for every $n > n_0^l$, the arrival rate is $\lambda^l P_{i_l}$, as $\sigma^l(S_{i_l}|n) = 1$ for all $n > n_0^l$. If $\lambda^l P_{i_l} \geq 1$, then the corresponding Markov chain would be unstable, and the IC constraint corresponding to i_l would be violated, which contradicts the fact that strong duality holds. Consequently, $\lambda^l P_{i_l} < 1$. Now, we can simply pick $n_0 = \max_{l \in [L]} n_0^l$ and $\epsilon = \min_{l \in [L]}(1 - \lambda^l P_{i_l})$. This completes the proof of Theorem 1.

7. Conclusion: discussion and future work

In this work, we tackle the problem of designing a congestion signaling mechanism where the server aims to maximize its throughput with utility-maximizing customers. We focus on two real-world features that fundamentally affect system throughput and customer utility: (i) customer heterogeneity with private types (i.e., private patience information), and (ii) demand variation. Our findings reveal that customer heterogeneity favors customer classes with high patience, enabling them to receive strictly positive expected utility, even though low-patience customers do not receive positive utility in expectation. Conversely, demand variation benefits the server by providing additional flexibility due to customers' lack of arrival rate information. We leverage a Bayesian persuasion framework and a C-MDP formulation of the problem to uncover the structural properties of the server's optimal congestion signaling mechanism. Our results show that customer heterogeneity can lead to a counter-intuitive laminar structure in the optimal policy.

We next discuss some extensions of our results and suggest directions for future research. Our framework for obtaining the structure of the optimal policy extends beyond the congestion signaling problem. Our methodology can be applied to a class of admission control problems with a linear reward function, where the optimal policy also exhibits a laminar structure. Another extension is the revenue maximization problem, where the server selects both a price and a signaling mechanism. Our results imply that for any fixed price, the server uses a recommendation-obedient signaling mechanism with ordered laminar policy, and so it holds under the optimal price also. More details are provided in Appendix A.1 and Appendix A.2.

In this work, we consider the case where the customer's utility function decays linearly with wait time. Some results can be extended to other classes of utility functions, such as ordered utility functions or monotonically decreasing utility functions. More details are provided in Appendix A.3.

Connection to learning: The C-MDP formulation of the optimal signaling problem opens up new avenues. For example, it allows for a flexible choice of objectives (throughput being a possible objective), expanding the scope of signal design considerations. Further, reinforcement learning algorithms can be employed to learn the optimal signaling mechanism, when the system parameters are unknown. Using a learning-based algorithm to derive the optimal policy would be an interesting direction for future research.

References

Akşin Z, Ata B, Emadi SM, Su CL (2017) Impact of delay announcements in call centers: An empirical approach. *Operations Research* 65(1):242–265.

Aldor-Noiman S, Feigin PD, Mandelbaum A (2009) Workload forecasting for a call center: Methodology and a case study .

Allon G, Bassamboo A, Gurvich I (2011) "we will be right with you": Managing customer expectations with vague promises and cheap talk. *Operations research* 59(6):1382–1394.

Altman E (2021) *Constrained Markov decision processes* (Routledge).

Anunrojwong J, Iyer K, Lingenbrink D (2024) Persuading risk-conscious agents: A geometric approach. *Operations Research* 72(1):151–166.

Anunrojwong J, Iyer K, Manshadi V (2023) Information design for congested social services: Optimal need-based persuasion. *Management Science* 69(7):3778–3796.

Arieli I, Babichenko Y (2019) Private bayesian persuasion. *Journal of Economic Theory* 182:185–217.

Armony M, Maglaras C (2004) Contact centers with a call-back option and real-time delay information. *Operations research* 52(4):527–545.

Armony M, Shimkin N, Whitt W (2009) The impact of delay announcements in many-server queues with abandonment. *Operations Research* 57(1):66–81.

Avramidis AN, Deslauriers A, L'Ecuyer P (2004) Modeling daily arrivals to a telephone call center. *Management Science* 50(7):896–908.

Bassamboo A, Randhawa RS, Zeevi A (2010) Capacity sizing under parameter uncertainty: Safety staffing principles revisited. *Management Science* 56(10):1668–1686.

Bergemann D, Morris S (2016a) Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics* 11(2):487–522.

Bergemann D, Morris S (2016b) Information design, bayesian persuasion, and bayes correlated equilibrium. *American Economic Review* 106(5):586–591.

Bergemann D, Morris S (2019) Information design: A unified perspective. *Journal of Economic Literature* 57(1):44–95.

Candogan O (2020) Information design in operations. *Pushing the Boundaries: Frontiers in Impactful OR/OM Research*, 176–201 (INFORMS).

Candogan O, Strack P (2023) Optimal disclosure of information to privately informed agents. *Theoretical Economics* 18(3):1225–1269.

Chen H, Frank MZ (2001) State dependent pricing with a queue. *Iie Transactions* 33(10):847–860.

Chen Y, Dong J, Wang Z (2021) A primal-dual approach to constrained markov decision processes. *arXiv preprint arXiv:2101.10895* .

Crawford VP, Sobel J (1982) Strategic information transmission. *Econometrica: Journal of the Econometric Society* 1431–1451.

Dong J, Yom-Tov E, Yom-Tov GB (2019) The impact of delay announcements on hospital network coordination and waiting times. *Management Science* 65(5):1969–1994.

Dworczak P, Martini G (2019) The simple economics of optimal persuasion. *Journal of Political Economy* 127(5):1993–2048.

Economou A (2021) The impact of information structure on strategic behavior in queueing systems. *Queueing theory* 2:137–169.

Edelson NM, Hilderbrand DK (1975) Congestion tolls for poisson queuing processes. *Econometrica: Journal of the Econometric Society* 81–92.

Efroni Y, Mannor S, Pirotta M (2020) Exploration-exploitation in constrained mdps. *arXiv preprint arXiv:2003.02189* .

Farrell J, Rabin M (1996) Cheap talk. *Journal of Economic perspectives* 10(3):103–118.

Guo P, Zipkin P (2007) Analysis and comparison of queues with different levels of delay information. *Management Science* 53(6):962–970.

Guo X, Hernández-Lerma O, Guo X, Hernández-Lerma O (2009) *Continuous-time Markov decision processes* (Springer).

Hassin R (1986) Consumer information in markets with random product quality: The case of queues and balking. *Econometrica: Journal of the Econometric Society* 1185–1195.

Hassin R (2016) *Rational queueing* (CRC press).

Hassin R, Haviv M (2003) *To queue or not to queue: Equilibrium behavior in queueing systems*, volume 59 (Springer Science & Business Media).

Hopenhayn H, Rogerson R (1993) Job turnover and policy evaluation: A general equilibrium analysis. *Journal of political Economy* 101(5):915–938.

Hopenhayn HA (1992) Entry, exit, and firm dynamics in long run equilibrium. *Econometrica: Journal of the Econometric Society* 1127–1150.

Ibrahim R (2018) Sharing delay information in service systems: a literature survey. *Queueing Systems* 89(1):49–79.

Ibrahim R, Ye H, L'Ecuyer P, Shen H (2016) Modeling and forecasting call center arrivals: A literature survey and a case study. *International Journal of Forecasting* 32(3):865–874.

Kamenica E, Gentzkow M (2011) Bayesian persuasion. *American Economic Review* 101(6):2590–2615.

Kartik N (2009) Strategic communication with lying costs. *The Review of Economic Studies* 76(4):1359–1395.

Khan N, Subramanian V (2023) A strong duality result for constrained pomdps with multiple cooperative agents. *arXiv preprint arXiv:2303.14932* .

Kolotilin A (2018) Optimal information disclosure: A linear programming approach. *Theoretical Economics* 13(2):607–635.

Kolotilin A, Mylovanov T, Zapechelnyuk A, Li M (2017) Persuasion of a privately informed receiver. *Econometrica* 85(6):1949–1964.

Kostami V, Ward AR (2009) Managing service systems with an offline waiting option and customer abandonment. *Manufacturing & Service Operations Management* 11(4):644–656.

Krishna V (2009) *Auction theory* (Academic press).

Krishna V, Morgan J (2001) A model of expertise. *The Quarterly Journal of Economics* 116(2):747–775.

Lingenbrink D, Iyer K (2019) Optimal signaling mechanisms in unobservable queues. *Operations research* 67(5):1397–1416.

Naor P (1969) The regulation of queue size by levying tolls. *Econometrica: journal of the Econometric Society* 15–24.

Nash P, Anderson EJ (1987) Linear programming in infinite-dimensional spaces: theory and applications. *(No Title)* .

Nirenberg S, Daw A, Pender J (2018) The impact of queue length rounding and delayed app information on disney world queues. *2018 Winter Simulation Conference (WSC)*, 3849–3860 (IEEE).

Puterman ML (2014) *Markov decision processes: discrete stochastic dynamic programming* (John Wiley & Sons).

Rayo L, Segal I (2010) Optimal information disclosure. *Journal of political Economy* 118(5):949–987.

Roy A, Borkar V, Karandikar A, Chaporkar P (2021) Online reinforcement learning of optimal threshold policies for markov decision processes. *IEEE Transactions on Automatic Control* 67(7):3722–3729.

Sutton RS, Barto AG (2018) *Reinforcement learning: An introduction* (MIT press).

Walker K, Potter E, Hwang I, Dwyer T, Egerton-Warburton D, Joe K, Hutton J, Freeman S, Flynn D (2022) Visualising emergency department wait times; rapid iterative testing to determine patient preferences for displays. *medRxiv* 2022–03.

Wei H, Liu X, Ying L (2024) Safe reinforcement learning with instantaneous constraints: The role of aggressive exploration. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 21708–21716.

Whitt W (1999) Improving service by informing customers about anticipated delays. *Management science* 45(2):192–207.

Willem M (2012) *Minimax theorems*, volume 24 (Springer Science & Business Media).

Wolff RW (1982) Poisson arrivals see time averages. *Operations research* 30(2):223–231.

Yu Q, Allon G, Bassamboo A (2017) How do delay announcements shape customer behavior? an empirical study. *Management Science* 63(1):1–20.

Appendix A: Additional result on admission control and revenue maximization

In this section, we provide extensions of the results presented in this work. For simplicity, in this section, we consider that the arrival is fixed, i.e., $\Lambda = \lambda$ with probability 1. We start with considering a class of admission control problems.

A.1. Admission control with linearly decaying reward

In this section, we consider a class of admission control problem where the reward function under any action decays linearly with respect to the queue length. Using similar notations as in Section 2, we assume that server has $K+1$ possible actions given by $\mathcal{S} = \{S_1, S_2, \dots, S_{K+1}\}$. Each action S_i allows the admission of a cumulative P_i fraction of customers, i.e., the arrival rate under action S_i is λP_i. Without loss of generality, we assume that $P_1 \geq P_2 \geq \cdots \geq P_{K+1}$. Upon taking action S_i, the reward received by the server is $H_i(n)$ when the system's current state (given by queue length) is n. We consider the function $H_i(\cdot)$ to be linear and monotonically decreasing for all $i \in [K+1]$, i.e., $H_i(n) := \alpha_i - \gamma_i n$ where $\gamma_i \geq 0$. We assume that for action S_{K+1}, $P_{K+1} = 0$ and $\alpha_{K+1} = \gamma_{K+1} = 0$, i.e., the arrival rate and reward are both zero under the action S_{K+1}, and for all other actions, i.e., S_i for $i \in [K]$, we have $\gamma_i > 0$. The server's objective is to maximize the overall reward. We show that for the admission control problem, the optimal policy has a laminar structure.

Proposition 6 (Admission control)**.** *Consider the optimization problem*

$$\textit{Adm-Con OPT:} \quad \max_{\boldsymbol{\sigma}} \; H(\boldsymbol{\sigma}), \; \text{where } H(\boldsymbol{\sigma}) := \mathbb{E}_{\boldsymbol{\sigma}}\Big[\sum_{i=1}^{K} \xi(S_i) H_i(W)\Big], \tag{14}$$

where $H_i(\cdot)$ is given by $H_i(n) := \alpha_i - \gamma_i n$ where $\gamma_i \geq 0$. Further, under action S_{K+1}, $P_{K+1} = 0$ and $\alpha_{K+1} = \gamma_{K+1} = 0$ and $\gamma_i > 0$ for $i \in [K]$. Then, there exists an optimal policy $\boldsymbol{\sigma}^*$ that is ordered laminar.

Proposition 6 demonstrates that the laminar structure of the optimal policy is prevalent across a wider range of problems beyond throughput optimization. The proof of Proposition 6 leverages similar arguments to those used in the proof of Theorem 1, illustrating that the methodology developed in this work is applicable beyond just the throughput optimization problem (4).

A.2. Revenue maximization

In this work, we consider the problem where the service provider wants to maximize the system throughput, equivalent to maximizing the revenue when every customer pays a fixed price r. Instead of a fixed price, the service provider might use the price level as a lever to optimize revenue; that is to say, the server first decides a fixed price r and then chooses a throughput maximizing signaling mechanism compatible with the customer's incentives. The customer's utility also depreciates with an increase in the price level r, and the customer of type-i receives the utility $h_i(n) - r = c_i - r - b_i n$ upon entry. Without loss of generality, we assume that $\frac{c_i - r}{b_i}$ is strictly increasing with respect to i, as otherwise, we can relabel the indices.

The reasoning behind considering a single price level r is that, in this work, we assume that the service level of each customer is the same; that is, the service requirement of each of the customers is exponentially distributed with mean 1. And as the service requirement is the same, the cost should be the same.

Corollary 3. *For a given price level r, there exists an optimal signaling mechanism $\mathbf{\Sigma}^{(r)} = (\mathcal{S}^{[K]}, \boldsymbol{\sigma}^{(r)})$ such that the optimal policy $\boldsymbol{\sigma}^{(r)}$ is the solution of the following C-MDP problem.*

$$\textit{Revenue C-MDP: } \max_{\boldsymbol{\sigma}} \mathbb{E}_{\boldsymbol{\sigma}}\Big[\sum_{i=1}^{K} P_i \xi(S_i)\Big]$$
$$\textit{s.t. } \mathbb{E}_{\boldsymbol{\sigma}}\big[\xi(S_i)(h_i(W)-r)\big] \geq 0, \ \forall i \in [K], \tag{15}$$

where $P_i := \sum_{j\geq i} p_j$. Further, the optimal solution $\boldsymbol{\sigma}^{(r)}$ is ordered laminar. Suppose Re(r) is the revenue generated under the policy $\boldsymbol{\sigma}^{(r)}$ under price level r. Then, the optimal price level is $r^* = \arg\max_{r\geq 0}$ Re(r).

Corollary 3 presents the revenue maximization problem where the server uses both the price level and the signaling mechanism as levers to optimize the revenue. The results in Corollary 3 follow from the results in Proposition 3 and Theorem 1. For a given price level r, Revenue C-MDP presents the C-MDP formulation for finding the optimal signaling mechanism $\mathbf{\Sigma}^{(r)} = (\mathcal{S}^{[K]}, \boldsymbol{\sigma}^{(r)})$. In the Revenue C-MDP in (15), we have only included the first IC constraint, i.e. $\mathbb{E}_{\boldsymbol{\sigma}}\big[\xi(S_i)(h_i(W)-r)\big] \geq 0$ for all $i \in [K]$, as from Proposition 3, we have that the IC constraint $\mathbb{E}_{\boldsymbol{\sigma}}\big[\xi(S_{i+1})(h_i(W)-r)\big] \leq 0$ is non-binding for all $i \in [K+1]$.

We note that the function Re(r) need not be continuous, and we observe that, in many cases, the function Re(r) is discontinuous at the values $\{c_i\}_{i\in[K]}$. The reason behind this is that when the price level r exceeds the type-i customer's value of the service c_i (that is $r > c_i$), the type-i customer has no incentive to enter the queue, even if the queue is empty. However, when the r is just below c_i, the type-i is incentivized to join the queue if it is empty. This phenomenon sometimes leads to a significant difference between the maximum revenue when the price level r is just below c_i and just above c_i, which leads to discontinuities in the function Re(r).

A.3. General utility function

In this work, we consider the case where the customer's utility function decays linearly with wait time. Some results can be extended to other classes of utility functions.

Ordered utility function: The utility functions h_i are considered *ordered* if there exist positive constants b_i's such that $\frac{1}{b_i}h_i(n)$'s maintain a strict order. Specifically, without loss of generality, this order is given by $\frac{1}{b_1}h_1(n) < \frac{1}{b_2}h_2(n) < \cdots < \frac{1}{b_K}h_K(n)$, for all $n \geq 0$. In the context of heterogeneous customers with private types, the results presented in Proposition 1 and Proposition 2 hold when the utility functions h_i are ordered. However, it is important to note that the existence of an optimal policy with a laminar structure is not guaranteed in this case, meaning that the result in Theorem 1 may not necessarily apply.

For *general* utility functions, not necessarily ordered, it can be shown that the server only needs to use at most 2^K signals. These signals correspond to the binary vectors in $\{0,1\}^K$, representing the possible responses ("Enter" or "Leave") of K different customer types. This allows the results in Proposition 1 and Proposition 2 to be extended to a broader class of utility functions.

Monotonically decreasing utility function with homogeneous customers: When customers are homogeneous, the server can use binary signals ("Enter" or "Leave"). Further, we hypothesize that if the utility function is monotonically decreasing, an optimal policy with a threshold structure exists, i.e., the result in Proposition 4 extend to any monotonically decreasing utility function $h(\cdot)$. Additionally, the result in Proposition 5 holds even for non-linear utility functions $h(\cdot)$, i.e., the server benefits from withholding the realization of the arrival rate.

Although some results in this work can be extended to a broader class of utility functions (as mentioned above), the structure of the optimal policy remains unknown, and characterizing this structure is a direction for future research.

Appendix B: Details on monotone and laminar policy structure

B.1. Equivalent definition of monotone policy

Lemma 5. *Consider a policy $\boldsymbol{\sigma} : \mathbb{Z}_+ \to \Omega(\mathcal{S}^{[K]})$. It is monotone if and only if the following holds: There exists an ordering $k : [K+1] \to [K+1]$ such that, for all $i, j \in [K+1]$, $\sigma(S_{k(i)}|n) > 0$ implies that for all $n' > n$, $\sigma(S_{k(j)}|n') = 0$ for all $j < i$.*

The monotone policy characterized in Lemma 5 is more general than the one discussed in Section 4.1. Specifically, the policy $\boldsymbol{\sigma}$ over the patience cutoff signals is considered ordered monotone (in line with the Definition 7) when $k(i) = i$. To appreciate an implication of Lemma 5, consider a case with two signals (i.e., $K = 1$), where $\sigma(S_1|n) = \sigma(S_2|n) = 1/2$ for all n, meaning that the server is providing signal S_1 with probability $1/2$ irrespective of the queue length. Such a policy is equivalent to a *no information* policy, as the posterior distribution of the queue length given the signal remains the same. Even though this policy has a "threshold" structure in terms of the probability with which signal S_1 is provided, the policy is not monotone. The proof of Lemma 5 is provided below.

B.1.1. Proof of Lemma 5 For ease of exposition, suppose statement $C.1$ is given as follows.

$C.1$*: There exists an ordering $k : [K+1] \to [K+1]$ such that, for all $i \in [K+1]$, $\sigma(S_{k(i)}|n) > 0$ implies that for all $n' > n$, $\sigma(S_{k(j)}|n') = 0$ for all $j < i$.*

Part 1 [$\boldsymbol{\sigma}$ is monotone implies $C.1$]:

- As $\boldsymbol{\sigma}$ is a monotone policy, there exists a monotone partition $\{\mathcal{T}_i\}_{i\in[K+1]}$ with $\mathcal{T}_i = [x_{i-1}, x_i)$ and $0 = x_0 \le x_1 \le x_2 \le \cdots \le x_K \le x_{K+1} = \infty$, and an ordering $k(\cdot)$, such that $\sigma(S_{k(i)}|n) = \mathbb{P}(n + X \in \mathcal{T}_i)$. For ease of notation, in this proof, we relabel the signals by using the notation $\tilde{S}_i = S_{k(i)}$ for all $i \in [K+1]$.

  As $\mathcal{T}_i = [x_{i-1}, x_i)$, we have $\sigma(\tilde{S}_i|n) = \mathbb{P}(n+X \in [x_{i-1}, x_i))$, and so $\sigma(\tilde{S}_i|n) > 0$ only if $n \ge \lfloor x_{i-1} \rfloor$. Consider $j < i$ and $n' > n$. We have $n' \ge n + 1 \ge \lfloor x_{i-1} \rfloor + 1 \ge \lfloor x_j \rfloor + 1$, where the last inequality use $x_{i-1} \ge x_j$ for $j < i$. As such, $\sigma(\tilde{S}_j|n') = \mathbb{P}(n' + X \in [x_{j-1}, x_j)) = 0$. As the argument holds for any $j < i$ and $n' > n$, we get that statement $C.1$ holds. □

Part 2 [$C.1$ implies $\boldsymbol{\sigma}$ is monotone]:

- To show this, we need to construct a monotone partition $\{\mathcal{T}_i\}_{i\in[K+1]}$ such that $\boldsymbol{\sigma}$ can be represented using $\{\mathcal{T}_i\}_{i\in[K+1]}$. Let $x_0 = 0$, and for any $i \in [K]$, choose x_i as $x_i = x_{i-1} + \sum_{n=0}^{\infty} \sigma(\tilde{S}_i|n)$, and finally, $x_{K+1} = \infty$. Define the set $\mathcal{T}_i = [x_{i-1}, x_i)$ for all $i \in [K+1]$. It is easy to see that $0 = x_0 \le x_1 \le x_2 \le \cdots \le x_K \le x_{K+1} = \infty$ and so $\{\mathcal{T}_i\}_{i\in[K+1]}$ is a monotone partition.

  Consider a given $i \in [K+1]$. We only need to show that $\sigma(\tilde{S}_i|n) = \mathbb{P}(n + X \in \mathcal{T}_i)$. Suppose $\overline{n}_i = \max\{n : \exists j \le i \text{ s.t. } \sigma(\tilde{S}_j|n) > 0\}$. Note that, by assumption that $C.1$ holds true, $\sigma(\tilde{S}_j|n) = 0$ for all $j > i$ and $n < \overline{n}_i$. As such, we have that $\sum_{j\le i} \sigma(\tilde{S}_j|n) = 1$ for all $0 \le n < \overline{n}_i$, and so,

$$x_i = \sum_{n=0}^{\infty} \sum_{j\le i} \sigma(\tilde{S}_j|n) = \sum_{n=0}^{\overline{n}_i} \sum_{j\le i} \sigma(\tilde{S}_j|n) = \overline{n}_i + \sum_{j\le i} \sigma(\tilde{S}_j|\overline{n}_i).$$

  This implies that for any n,

$$\mathbb{P}(n + X \in \cup_{j\le i}\mathcal{T}_j) = \mathbb{P}(n + X < x_i) = \begin{cases} 1 & \text{for } n < \overline{n}_i, \\ \sum_{j\le i} \sigma(\tilde{S}_j|\overline{n}_i) & \text{for } n = \overline{n}_i, \\ 0 & \text{for } n > \overline{n}_i. \end{cases}$$
$$= \sum_{j\le i} \sigma(\tilde{S}_j|n),$$

  where the last equality uses the definition of $\overline{n}_i$ and $\sum_{j\le i} \sigma(\tilde{S}_j|n) = 1$ for all $0 \le n < \overline{n}_i$. Now, as the sets $\{\mathcal{T}_i\}_{i\in[K+1]}$ are disjoint,

$$\begin{aligned}\mathbb{P}(n + X \in \mathcal{T}_i) &= \mathbb{P}(n + X \in \cup_{j\le i}\mathcal{T}_j) - \mathbb{P}(n + X \in \cup_{j\le i-1}\mathcal{T}_j) \\ &= \sum_{j\le i} \sigma(\tilde{S}_j|n) - \sum_{j\le i-1} \sigma(\tilde{S}_j|n) \\ &= \sigma(\tilde{S}_i|n).\end{aligned}$$

  As the above argument holds for any $i \in [K+1]$, it shows that $\boldsymbol{\sigma}$ is a monotone policy with respect to the monotone partition $\{\mathcal{T}_i\}_{i\in[K+1]}$. This completes the proof. □

B.2. Proof of Lemma 1

For ease of exposition, suppose statement $C.2$ is given as follows.

$C.2$: For any $i, j \in [K+1]$, $\sigma(S_i|n) > 0$ for $n = \underline{n}_i$ and $n = \overline{n}_i$, where $\underline{n}_i \le \bar{n}_i$, implies that $\sigma(S_j|n) = 0$ for all $j < i$ and $\underline{n}_i < n < \overline{n}_i$.

Part 1 [$\boldsymbol{\sigma}$ is ordered laminar implies $C.2$]:

Suppose $\{\mathcal{T}_i\}_{i\in[K+1]}$ is a ordered laminar partition with respect to the ordered laminar interval family $\{\mathcal{I}_i\}_{i\in[K+1]}$. We have that, for any $i, j \in [K+1]$, with $j < i$, either $\mathcal{I}_i \cap \mathcal{I}_j = \emptyset$ or $\mathcal{I}_i \subseteq \mathcal{I}_j$. And for any $j \in [K+1]$, $\mathcal{T}_j = \mathcal{I}_j \setminus \cup_{i>j} \mathcal{I}_i$.

Now, suppose $\boldsymbol{\sigma}$ is an ordered laminar policy with respect to $\{\mathcal{T}_i\}_{i\in[K+1]}$. Then, for any $i \in [K+1]$, $\sigma(S_i|n) = \mathbb{P}(n+X \in \mathcal{T}_i)$. This implies that

$$\sum_{j\geq i} \sigma(S_j|n) = \sum_{j\geq i} \mathbb{P}(n+X\in\mathcal{T}_j) \overset{(a)}{=} \mathbb{P}(n+X\in\cup_{j\geq i}\mathcal{I}_j) \geq \mathbb{P}(n+X\in\mathcal{I}_i),$$

where (a) follows using $\mathcal{T}_j = \mathcal{I}_j \setminus \cup_{i>j}\mathcal{I}_i$.

Now, if there exists $\underline{n}_i$ and $\overline{n}_i$ such that $\sigma(S_i|n)>0$, then $\mathbb{P}(n+X\in\mathcal{I}_i)>0$ at $n=\underline{n}_i$ and $n=\overline{n}_i$. And as $\mathcal{I}_i$ is an interval, it implies that $[\underline{n}_i+1,\overline{n}_i)\subseteq\mathcal{I}_i$. This in turn implies that $\mathbb{P}(n+X\in\mathcal{I}_i)=1$ for any $\underline{n}_i<n<\overline{n}_i$. This finally gives us that, for any $j'<i$ and $\underline{n}_i<n<\overline{n}_i$,

$$\sigma(S_{j'}|n) \leq 1-\sum_{j\geq i}\sigma(S_j|n) \leq 1-\mathbb{P}(n+X\in\mathcal{I}_i)=0.$$

This proves that the statement $C.2$ holds.

Part 2 [$C.2$ implies $\boldsymbol{\sigma}$ is ordered laminar]:

We use the following iterative steps to create an ordered laminar interval family. Without loss of generality, we assume that $\sum_{n\geq 0}\sigma(S_i|n)>0$ for all $i\in[K+1]$.

1. In the first step, consider $i=K+1$ and let $\underline{N}_{K+1}=\min\{n:\sigma(S_{K+1}|n)>0\}$ and $\overline{N}_{K+1}=\max\{n:\sigma(S_{K+1}|n)>0\}$ ($\overline{N}_{K+1}=\infty$ is allowed). Now, the set $\mathcal{I}_{K+1}=[\underline{x}_{K+1},\overline{x}_{K+1})$, where

$$\underline{x}_{K+1} = \underline{N}_{K+1}+1-\sigma(S_{K+1}|\underline{N}_{K+1}), \quad \text{and}$$
$$\overline{x}_{K+1} = \begin{cases} \overline{N}_{K+1}+\sigma(S_{K+1}|\overline{N}_{K+1}) & \text{if } \overline{N}_{K+1}>\underline{N}_{K+1}, \\ \underline{N}_{K+1}+1 & \text{if } \overline{N}_{K+1}=\underline{N}_{K+1}. \end{cases}$$

   For any $n\geq 0$, if $[n,n+1)\cap\left[\mathcal{I}_{K+1}\right]^c\neq\emptyset$, we define

$$\texttt{Lower}_K(n)=\inf\{x:x\geq n \text{ and } x\notin\mathcal{I}_{K+1}\},$$
$$\texttt{Upper}_K(n)=\sup\{x:x\leq n+1 \text{ and } x\notin\mathcal{I}_{K+1}\}.$$

2. Consider a general $i\in[K]$ and let $\underline{N}_i=\min\{n:\sigma(S_i|n)>0\}$ and $\overline{N}_i=\max\{n:\sigma(S_i|n)>0\}$. Now, the set $\mathcal{I}_i=[\underline{x}_i,\overline{x}_i)$, where

$$\underline{x}_i=\texttt{Upper}_i(\underline{N}_i)-\sigma(S_i|\underline{N}_i), \quad \text{and} \quad \overline{x}_i=\begin{cases}\texttt{Lower}_i(\overline{N}_i)+\sigma(S_i|\overline{N}_i) & \text{if } \overline{N}_i>\underline{N}_i, \\ \texttt{Upper}_i(\underline{N}_i) & \text{if } \overline{N}_i=\underline{N}_i,\end{cases}$$

   For any $n\geq 0$, if $[n,n+1)\cap\left[\cup_{j\geq i}\mathcal{I}_j\right]^c\neq\emptyset$, we define

$$\texttt{Lower}_{i-1}(n)=\inf\{x:x\geq n \text{ and } x\notin\cup_{j\geq i}\mathcal{I}_j\},$$
$$\texttt{Upper}_{i-1}(n)=\sup\{x:x\leq n+1 \text{ and } x\notin\cup_{j\geq i}\mathcal{I}_j\}.$$

   Note that the quantities $\texttt{Lower}_i(n)$ and $\texttt{Upper}_i(n)$ might not exist for some n, however, we only need to define the $\texttt{Lower}_i(n)$ and $\texttt{Upper}_i(n)$ when $[n,n+1)\cap\left[\cup_{j\geq i}\mathcal{I}_j\right]^c\neq\emptyset$.

Claim 1. *For all $n \geq 0$, the following holds.*

(i) For any $i \in [K+1]$, $\mathbb{P}(n+X \in \cup_{j\geq i}\mathcal{I}_j) = \sum_{j\geq i}\sigma(S_j|n)$.

(ii) For any $i \in [K+1]$, the set $[n,n+1)\backslash\big[\cup_{j\geq i}\mathcal{I}_j\big]$ is either empty or

$$[n,n+1)\backslash\big[\cup_{j\geq i}\mathcal{I}_j\big] = \big[\texttt{Lower}_{i-1}(n), \texttt{Upper}_{i-1}(n)\big)$$

(iii) The set of subsets $\{\mathcal{I}_j\}_{j\geq i}$ forms an ordered laminar interval family.

We prove the claim using an induction based argument.

Consider that $i = K+1$: Then,

$$\begin{aligned}
\mathbb{P}(n+X \in \mathcal{I}_{K+1}) &= \mathbb{P}(n+X < [\underline{x}_{K+1}, \overline{x}_{K+1})) \\
&= \begin{cases} 1 & \text{for } \underline{N}_{K+1} < n < \overline{N}_{K+1}, \\ \sigma(S_{K+1}|n) & \text{for } n = \overline{N}_{K+1}, \text{ or } n = \underline{N}_{K+1}, \\ 0 & \text{otherwise.} \end{cases} \\
&\overset{(a)}{=} \sigma(S_{K+1}|n),
\end{aligned}$$

where (a) holds because $\sigma(S_i|n) = 0$ for all $i \in [K]$ and $\underline{N}_{K+1} < n < \overline{N}_{K+1}$ by using statement $C.2$, and so $\sigma(S_{K+1}|n) = 1$ for $\underline{N}_{K+1} < n < \overline{N}_{K+1}$. Note that the above argument holds even when $\overline{N}_{K+1} = \underline{N}_{K+1}$. This shows that part (i) of the claim holds. Now, part (ii) and (iii) follows simply as $\mathcal{I}_{K+1}$ is an interval.

Consider $i \in [K]$ and suppose the claim holds for all $j > i$ and $j \in [K+1]$: We show that the claim holds for i. As the claim holds for all $j > i$, it implies that for any n, by using part (i),

$$\mathbb{P}\big(X \in [n,n+1)\backslash\big[\cup_{j>i}\mathcal{I}_j\big]\big) = 1 - \mathbb{P}\big(n+X \in \cup_{j>i}\mathcal{I}_j\big) = 1 - \sum_{j>i}\sigma(S_j|n).$$

As such, $[n,n+1)\backslash\big[\cup_{j>i}\mathcal{I}_j\big] = \emptyset$ only if $\sum_{j>i}\sigma(S_j|n) = 1$. Now, by using part (ii), we also get

$$\mathbb{P}\big(X \in [n,n+1)\backslash\big[\cup_{j>i}\mathcal{I}_j\big]\big) = 1 - \sum_{j>i}\sigma(S_j|n) = \texttt{Upper}_i(n) - \texttt{Lower}_i(n), \tag{16}$$

where the last equality holds whenever $[n,n+1)\backslash\big[\cup_{j>i}\mathcal{I}_j\big] \neq \emptyset$.

This implies that at $n = \underline{N}_i$, as $\sigma(S_i|n) > 0$, we have $[n,n+1)\backslash\big[\cup_{j>i}\mathcal{I}_j\big] \neq \emptyset$. The same argument holds for $n = \overline{N}_i$. As such, the terms $\texttt{Lower}_i(n)$ and $\texttt{Upper}_i(n)$ are well defined at $n = \underline{N}_i$ and $n = \overline{N}_i$. And

$$\texttt{Upper}_i(n) - \texttt{Lower}_i(n) = 1 - \sum_{j>i}\sigma(S_j|n) \geq \sigma(S_i|n), \quad \text{for } n = \underline{N}_i \text{ and } n = \overline{N}_i. \tag{17}$$

Next, we show that part (i) and (ii) are satisfied for the given i and for all $n \geq 0$.

1. **Case 1:** Suppose $n < \underline{N}_i$ or $n > \overline{N}_i$. For all such n, we have $[n, n+1) \cap \mathcal{I}_i = \emptyset$ and $\sigma(S_i|n) = 0$. This implies that

$$[n, n+1) \backslash \big[\cup_{j \geq i} \mathcal{I}_j \big] = [n, n+1) \backslash \big[\cup_{j > i} \mathcal{I}_j \big] = [\texttt{Lower}_i(n), \texttt{Upper}_i(n)),$$

and so, we choose $\texttt{Lower}_{i-1}(n) = \texttt{Lower}_i(n)$ and $\texttt{Upper}_{i-1}(n) = \texttt{Upper}_i(n)$. Again by using $[n, n+1) \cap \mathcal{I}_i = \emptyset$,

$$\mathbb{P}(n + X \in \cup_{j \geq i} \mathcal{I}_j) = \mathbb{P}(n + X \in \cup_{j > i} \mathcal{I}_j) \overset{(a)}{=} \sum_{j > i} \sigma(S_j|n) \overset{(b)}{=} \sum_{j \geq i} \sigma(S_j|n),$$

where (a) holds by using part (i) for $j > i$; and (b) holds because $\sigma(S_i|n) = 0$ for $n < \underline{N}_i$ or $n > \overline{N}_i$.

2. **Case 2:** Suppose $n = \underline{N}_i = \overline{N}_i$. In this case, we have $\mathcal{I}_i = [\underline{x}_i, \overline{x}_i) = \big[\texttt{Upper}_i(n) - \sigma(S_i|n), \texttt{Upper}_i(n)\big)$. Then,

$$\begin{aligned} &[n, n+1) \backslash \big[\cup_{j \geq i} \mathcal{I}_j \big] \\ &= [\texttt{Lower}_i(n), \texttt{Upper}_i(n)) \backslash \big[\texttt{Upper}_i(n) - \sigma(S_i|n), \texttt{Upper}_i(n)\big) \\ &= \big[\texttt{Lower}_i(n), \texttt{Upper}_i(n) - \sigma(S_i|n)\big) \\ &= \begin{cases} \emptyset & \text{if } \texttt{Lower}_i(n) = \texttt{Upper}_i(n) - \sigma(S_i|n) \\ \big[\texttt{Lower}_{i-1}(n), \texttt{Upper}_{i-1}(n)\big) & \text{otherwise,} \end{cases} \end{aligned}$$

where the last equality uses $\texttt{Upper}_{i-1}(n) = \texttt{Upper}_i(n) - \sigma(S_i|n)$ when $n = \underline{N}_i = \overline{N}_i$. Also,

$$\begin{aligned} \mathbb{P}(n + X \in \cup_{j \geq i} \mathcal{I}_j) &= \mathbb{P}(n + X \in \cup_{j > i} \mathcal{I}_j) \\ &\quad + \mathbb{P}\big(n + X \in [\texttt{Upper}_i(n) - \sigma(S_i|n), \texttt{Upper}_i(n))\big) \\ &\overset{(a)}{=} \sum_{j > i} \sigma(S_j|n) + \sigma(S_i|n) = \sum_{j \geq i} \sigma(S_j|n), \end{aligned}$$

where (a) uses that the claim holds for all $j > i$.

3. **Case 3:** Suppose $n = \overline{N}_i > \underline{N}_i$. In this case, we have $\mathcal{I}_i = [\underline{x}_i, \overline{x}_i)$, where $\overline{x}_i = \texttt{Lower}_i(n) + \sigma(S_i|n)$, at $n = \overline{N}_i$. Then,

$$\begin{aligned} &[n, n+1) \backslash \big[\cup_{j \geq i} \mathcal{I}_j \big] \\ &= [\texttt{Lower}_i(n), \texttt{Upper}_i(n)) \backslash \big[\texttt{Lower}_i(n) + \sigma(S_i|n), \texttt{Upper}_i(n)\big) \\ &= \big[\texttt{Lower}_i(n) + \sigma(S_i|n), \texttt{Upper}_i(n)\big) \\ &= \begin{cases} \emptyset & \text{if } \texttt{Lower}_i(n) + \sigma(S_i|n) = \texttt{Upper}_i(n) \\ \big[\texttt{Lower}_{i-1}(n), \texttt{Upper}_{i-1}(n)\big) & \text{otherwise,} \end{cases} \end{aligned}$$

where the last equality uses $\texttt{Lower}_{i-1}(n) = \texttt{Lower}_i(n) + \sigma(S_i|n)$ when $n = \overline{N}_i > \underline{N}_i$.

Also, we have

$$\begin{aligned}\mathbb{P}(n+X\in\cup_{j\geq i}\mathcal{I}_j) &= \mathbb{P}(n+X\in\cup_{j>i}\mathcal{I}_j)\\ &\quad+\mathbb{P}\big(n+X\in[\texttt{Lower}_i(n),\texttt{Lower}_i(n)+\sigma(S_i|n))\big)\\ &\overset{(a)}{=}\sum_{j>i}\sigma(S_j|n)+\sigma(S_i|n)=\sum_{j\geq i}\sigma(S_j|n),\end{aligned}$$

where (a) uses that the claim holds for all $j>i$.

4. **Case 4:** Suppose $n=\underline{N}_i<\overline{N}_i$. In this case, the arguments follows on similar lines as in Case 2 and Case 3 , by using $\underline{x}_i=\texttt{Upper}_i(n)-\sigma(S_i|n),$ at $n=\underline{N}_i$.
5. **Case 5:** In the last case, suppose $\underline{N}_i<n<\overline{N}_i$. In this case, as the statement $C.2$ holds, we have that for any $\underline{N}_i<n<\overline{N}_i$, $\sigma(S_j|n)=0$ for all $j<i$. This implies that for any $\underline{N}_i<n<\overline{N}_i$, $\sum_{j\geq i}\sigma(S_j|n)=1$. Now, as $[n,n+1)\subseteq\mathcal{I}_i$ for $\underline{N}_i<n<\overline{N}_i$, we have that $[n,n+1)\backslash\big[\cup_{j\geq i}\mathcal{I}_j\big]=\emptyset$. This also implies that

$$\mathbb{P}(n+X\in\cup_{j\geq i}\mathcal{I}_j)=1=\sum_{j\geq i}\sigma(S_j|n),\quad\text{for } \underline{N}_i<n<\overline{N}_i$$

From the above mentioned cases, we get that part (i) and (ii) holds for the given i, for any value of n. Finally, note that $\{\mathcal{I}_j\}_{j>i}$ is an ordered laminar interval family by assumption. Then, $\cup_{j>i}\mathcal{I}_j=\cup_{j>i}[\underline{x}_j,\overline{x}_i)$. Also, as $\underline{x}_i\in[\texttt{Lower}_i(n),\texttt{Upper}_i(n))$ at $n=\underline{N}_i$ and $\overline{x}_i\in[\texttt{Lower}_i(n),\texttt{Upper}_i(n))$ at $n=\overline{N}_i$, we have $\underline{x}_i\notin\cup_{j>i}\mathcal{I}_j$ and $\overline{x}_i\notin\cup_{j>i}\mathcal{I}_j$. This implies that for any $j>i$, either $[\underline{x}_j,\overline{x}_j)\cap[\underline{x}_i,\overline{x}_j)=\emptyset$ or $[\underline{x}_j,\overline{x}_j)\subset[\underline{x}_i,\overline{x}_j)$, and so $\{\mathcal{I}_j\}_{j\geq i}$ is also an ordered laminar interval family.

From the above cases, we have that the claim is satisfied at i, and so by the induction argument, the claim is true for all i. This completes the proof of the claim. □

From Claim 1, we have that $\{\mathcal{I}_i\}_{i\in[K+1]}$ is an ordered laminar interval family. Consider the corresponding laminar partition $\{\mathcal{T}_i\}_{i\in[K+1]}$. Then, by part (i) of Claim 1, for any i,

$$\begin{aligned}\sigma(S_i|n)&=\sum_{j\geq i}\sigma(S_j|n)-\sum_{j>i}\sigma(S_j|n)\\ &=\mathbb{P}(n+X\in\cup_{j\geq i}\mathcal{I}_j)-\mathbb{P}(n+X\in\cup_{j>i}\mathcal{I}_j)\\ &=\mathbb{P}\Big(n+X\in\cup_{j\geq i}\mathcal{I}_j\backslash\cup_{j>i}\mathcal{I}_j\Big)\\ &=\mathbb{P}\Big(n+X\in\mathcal{I}_i\backslash\cup_{j>i}\mathcal{I}_j\Big)=\mathbb{P}(n+X\in\mathcal{T}_i).\end{aligned}$$

This shows that $\boldsymbol{\sigma}$ is ordered laminar with respect to the partition $\{\mathcal{T}_i\}_{i\in[K+1]}$. This completes the proof.

Appendix C: Proof of Proposition 1 and Proposition 2

C.1. Proof of Proposition 1

Suppose $\boldsymbol{\Sigma}=(\mathcal{S},\boldsymbol{\sigma})$ is a given signaling mechanism and let $\boldsymbol{\nu}:\mathcal{S}\to[0,1]^K$ be a corresponding customer equilibrium. Note that, from Eq. (2), we have $\nu_i(S)>0$ only if $\mathbb{E}_{\boldsymbol{\Pi},\boldsymbol{\Sigma}}[\xi(S)h_i(W)]\geq 0$.

Claim 1: $\nu_i(S)>0$ implies that $\nu_{i^+}(S)=1$ for all $i^+>i$.

- The argument is as follows.

$$\begin{aligned}\mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}[\xi(S)h_i(W)] \geq 0 &\implies \mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}[\xi(S)W] \leq \frac{c_i}{b_i}\mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}[\xi(S)] \\ &\implies \mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}[\xi(S)W] < \frac{c_{i^+}}{b_{i^+}}\mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}[\xi(S)],\ \forall i^+ > i \\ &\implies \mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}[\xi(S)h_{i^+}(W)] > 0,\ \forall i^+ > i, \end{aligned} \tag{18}$$

where we use $h_i(n) = c_i - b_i n$ for all $i \in [K]$, and c_i/b_i follow an ascending order. And as $\mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}[\xi(S)h_{i^+}(W)] > 0$, we have $\nu_{i^+}(S) = 1$ for all $i^+ > i$. □

First, we prove that for any signaling mechanism $\mathbf{\Sigma}$ and a corresponding customer equilibrium $\boldsymbol{\nu}$, the server can design an equivalent signaling mechanism such that the resulting customer equilibrium is a pure strategy. For every $i \in [K]$, we define $\mathcal{S}_{(i)} \subseteq \mathcal{S}$ such that

$$\mathcal{S}_{(i)} = \{S \in \mathcal{S} : \nu_{i^-}(S) = 0 \ \forall\, i^- < i, 0 < \nu_i(S) \leq 1\},$$

and $\mathcal{S}_{(K+1)} = \{S \in \mathcal{S} : \nu_i(S) = 0 \ \forall\, i \in [K]\}$. From claim 1, we have that $\{\mathcal{S}_{(i)}\}_{i\in[K+1]}$ form a partition of $\mathcal{S}$. We construct a new signaling mechanism $\hat{\mathbf{\Sigma}} = (\mathcal{S} \times \mathcal{S}^{[K]}, \hat{\boldsymbol{\sigma}})$ such that, for any $i \in [K+1]$ and $S \in \mathcal{S}_{(i)}$,

$$\hat{\sigma}\big((S,S_i)|n,\lambda_l\big) = \nu_i(S)\sigma\big(S|n,\lambda_l\big) \text{ and } \hat{\sigma}\big((S,S_{i+1})|n,\lambda_l\big) = (1-\nu_i(S))\sigma\big(S|n,\lambda_l\big), \tag{19}$$

and for the remaining elements in $\mathcal{S} \times \mathcal{S}^{[K]}$, $\hat{\sigma}\big(\cdot|n,\lambda_l\big) = 0$. Under the signaling mechanism $\hat{\mathbf{\Sigma}}$, we consider the customer's strategy vector to be $\hat{\boldsymbol{\nu}}$ such that $\hat{\nu}\big((S,S_i)\big) = S_i$ and $\hat{\nu}\big((S,S_{i+1})\big) = S_{i+1}$.

<u>Claim 2: $\hat{\boldsymbol{\nu}}$ is a customer equilibrium under $\hat{\mathbf{\Sigma}}$.</u>

- We claim that the considered customer strategy vector $\hat{\boldsymbol{\nu}}$ is a customer equilibrium under the signaling mechanism $\hat{\mathbf{\Sigma}}$. First, note that the transition rates under the signaling mechanism $\hat{\mathbf{\Sigma}}$, assuming that the customer strategy vector is $\hat{\boldsymbol{\nu}}$, is the same as the transition rates under $\mathbf{\Sigma}$ with customer strategy vector $\boldsymbol{\nu}$: For the transition $(n,\lambda_l) \to (n+1,\lambda_l)$, by using the fact that $\{\mathcal{S}_{(i)}\}_{i\in[K+1]}$ is a partition of $\mathcal{S}$, the corresponding transition rate is

$$\begin{aligned}&\lambda_l \sum_{i=1}^{K+1}\sum_{S\in\mathcal{S}_{(i)}} \Big[\hat{\sigma}\big((S,S_i)|n,\lambda_l\big)\sum_{j=1}^{K}\hat{\nu}_j\big((S,S_i)\big)p_j + \hat{\sigma}\big((S,S_{i+1})|n,\lambda_l\big)\sum_{j=1}^{K}\hat{\nu}_j\big((S,S_{i+1})\big)p_j\Big] \\ &\overset{(a)}{=} \lambda_l \sum_{i=1}^{K+1}\sum_{S\in\mathcal{S}_{(i)}} \Big[\sigma\big(S|n,\lambda_l\big)\nu_i(S)\sum_{j=i}^{K}p_j + \sigma\big(S|n,\lambda_l\big)(1-\nu_i(S))\sum_{j=i+1}^{K}p_j\Big] \\ &= \lambda_l \sum_{i=1}^{K+1}\sum_{S\in\mathcal{S}_{(i)}} \sigma\big(S|n,\lambda_l\big)\Big[\nu_i(S)p_i + \sum_{j=i+1}^{K}p_j\Big] \\ &\overset{(b)}{=} \lambda_l \sum_{i=1}^{K+1}\sum_{S\in\mathcal{S}_{(i)}} \sigma\big(S|n,\lambda_l\big)\Big[\sum_{j=1}^{K}\nu_j(S)p_j\Big],\end{aligned}$$

where (a) follows by using Eq. (19) and, by assumption, $\hat{\nu}\big((S,S_i)\big)=S_i$ and $\hat{\nu}\big((S,S_{i+1})\big)=S_{i+1}$ for any $S\in\mathcal{S}_{(i)}$; and (b) follows by using the definition of $\mathcal{S}_{(i)}$. Thus, the transition rate for the transition $(n,\lambda_l)\to(n+1,\lambda_l)$ remains the same under $\hat{\mathbf{\Sigma}}$ and $\hat{\boldsymbol{\nu}}$. Similarly, the transition rate for the transition $(n,\lambda_l)\to(n-1,\lambda_l)$ also remains the same.

As the transition rates under $\hat{\mathbf{\Sigma}}$ and $\hat{\boldsymbol{\nu}}$ and those under $\mathbf{\Sigma}$ and $\boldsymbol{\nu}$ are the same, the steady-state distribution tuple $\hat{\mathbf{\Pi}}(\hat{\mathbf{\Sigma}},\hat{\boldsymbol{\nu}})$ is the same as the steady-state distribution tuple $\mathbf{\Pi}(\mathbf{\Sigma},\boldsymbol{\nu})$.

Next, for any $i\in[K]$ and $S\in\mathcal{S}_{(i)}$, and any function f, we have that

$$\begin{aligned}\mathbb{E}_{\hat{\mathbf{\Pi}},\hat{\mathbf{\Sigma}}}\big[\xi\big((S,S_i)\big)f(W)\big]&=\sum_{n=0}^{\infty}\sum_{l\in[L]}q_l\hat{\pi}_n^l\hat{\sigma}\big((S,S_i)|n,\lambda_l\big)f(n)\\&=\sum_{n=0}^{\infty}\sum_{l\in[L]}q_l\pi_n^l\sigma(S|n,\lambda_l)\nu_i(S)f(n)=\nu_i(S)\mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}\big[\xi(S)f(W)\big].\end{aligned}$$

This implies that for $i\in[K]$ and $S\in\mathcal{S}_{(i)}$, as $\nu_i(S)>0$,

$$\begin{aligned}\mathbb{E}_{\hat{\mathbf{\Pi}},\hat{\mathbf{\Sigma}}}\big[\xi\big((S,S_i)\big)h_i(W)\big]&=\nu_i(S)\mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}\big[\xi(S)h_i(W)\big]\geq 0\\\mathbb{E}_{\hat{\mathbf{\Pi}},\hat{\mathbf{\Sigma}}}\big[\xi\big((S,S_i)\big)h_{i^+}(W)\big]&=\nu_i(S)\mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}\big[\xi(S)h_{i^+}(W)\big]\overset{(a)}{>}0,\ \forall i^+>i,\\\mathbb{E}_{\hat{\mathbf{\Pi}},\hat{\mathbf{\Sigma}}}\big[\xi\big((S,S_i)\big)h_{i^-}(W)\big]&=\nu_i(S)\mathbb{E}_{\mathbf{\Pi},\mathbf{\Sigma}}\big[\xi(S)h_{i^-}(W)\big]\overset{(b)}{\leq}0,\ \forall i^->i,\end{aligned}$$

where (a) follows using the arguments in Eq. (18); and (b) uses that fact that $\mathbb{E}_{\boldsymbol{\sigma}}\big[\xi(S)h_{i^-}(W)\big]\leq 0$ for all $i^-<i$ as $\nu_{i^-}(S)=0$ for $i\in[K]$ and $S\in\mathcal{S}_{(i)}$. This implies that $\hat{\nu}\big((S,S_i)\big)=S_i$ satisfies the condition in Eq. (2). Using similar arguments, we have that $\hat{\nu}\big((S,S_{i+1})\big)=S_{i+1}$ also satisfies Eq. (2). Thus, $\hat{\boldsymbol{\nu}}$ is a customer equilibrium under $\hat{\mathbf{\Sigma}}$. □

Since the steady-state distribution tuple $\hat{\mathbf{\Pi}}(\hat{\mathbf{\Sigma}},\hat{\boldsymbol{\nu}})$ is the same as the steady-state distribution tuple $\mathbf{\Pi}(\mathbf{\Sigma},\boldsymbol{\nu})$, the signaling mechanism $\hat{\mathbf{\Sigma}}$ is equivalent to $\mathbf{\Sigma}$ (with customer equilibrium $\boldsymbol{\nu}$), where the customer equilibrium $\hat{\boldsymbol{\nu}}$ is a pure strategy.

In the next part, we assume that the server is implementing a signaling mechanism $\hat{\mathbf{\Sigma}}=(\hat{\mathcal{S}},\hat{\boldsymbol{\sigma}})$ under which the resulting customer equilibrium $\hat{\boldsymbol{\nu}}$ is a pure strategy. Further, as mentioned before $\hat{\nu}_i(S)>0$ implies that $\hat{\nu}_{i^+}(S)=1$ for all $i^+>i$. Thus, we have that $\hat{\boldsymbol{\nu}}(S)\in\{\mathbf{y}\in\{0,1\}^K:y_1\leq y_2\leq\cdots\leq y_K\}=\mathcal{S}^{[K]}$ for any $S\in\hat{\mathcal{S}}$. Now, suppose there are two signals two signals $S',S''\in\hat{\mathcal{S}}$ such that $\hat{\nu}_i(S')=\hat{\nu}_i(S'')$ for all $i\in[K]$, then the server can combine the two signals S' and S''. That is, the server can construct a signal $\tilde{S}=(S',S'')$. The new signaling scheme is given by $\tilde{\mathbf{\Sigma}}=(\tilde{\mathcal{S}},\tilde{\boldsymbol{\sigma}})$, where $\tilde{\mathcal{S}}=[\hat{\mathcal{S}}\backslash\{S',S''\}]\cup\{\tilde{S}\}$ and $\tilde{\sigma}(\tilde{S}|n,\lambda_l)=\hat{\sigma}(S'|n,\lambda_l)+\hat{\sigma}(S''|n,\lambda_l)$, and for all other signals, $\tilde{\sigma}(\cdot|\cdot,\cdot)$ remains the same. Suppose $\tilde{\boldsymbol{\nu}}$ is the resulting customer equilibrium under the signaling mechanism $\tilde{\mathbf{\Sigma}}$.

As the information provided by the signal $\tilde{S}$ is same as that by either S' or S'', the customer's response to the signal remains the same, i.e., $\tilde{\nu}_i(\tilde{S})=\hat{\nu}_i(S')=\hat{\nu}_i(S'')$ for all $i\in[K]$. Now, we can use similar arguments in the previous part to show that the transition probabilities of the underlying Markov chain remain the same, and so the signaling scheme $\tilde{\mathbf{\Sigma}}$ induces the same stationary distribution as $\hat{\mathbf{\Sigma}}$. This argument implies that the number of signals needed by the server is same as the number of unique customer strategy vectors $\hat{\boldsymbol{\nu}}(\cdot)$. And so, we can choose the set of signals to be $\mathcal{S}^{[K]}$ itself. In this case, for any plausible customer's response vector $S_i \in \mathcal{S}^{[K]}$, there is a unique signal $S \in \mathcal{S}^{[K]}$ such that $\tilde{\boldsymbol{\nu}}(S) = S_i$, as such, the server can just relabel the signals such that $\tilde{\boldsymbol{\nu}}(S_i) = S_i$ for all $S_i \in \mathcal{S}^{[K]}$.

C.2. Proof of Proposition 2

Using Proposition 1, we consider the signaling mechanism $\mathbf{\Sigma}$ to be such that $\mathbf{\Sigma} = (\mathcal{S}^{[K]}, \boldsymbol{\sigma})$ and obedience is the resulting customer equilibrium, i.e., $\boldsymbol{\nu}(S_i) = S_i$ for all $S_i \in \mathcal{S}^{[K]}$. In this case,

$$\begin{aligned}
\text{Th}(\mathbf{\Sigma}, \boldsymbol{\nu}) &= \sum_{l=1}^{L} q_l \Big[\sum_{n=0}^{\infty} \pi_n^l \sum_{i\in[K]} \lambda_l \sigma^l(S_i|n) \sum_{j=1}^{K} p_j \nu_j(S_i)\Big] \\
&\overset{(a)}{=} \sum_{l=1}^{L} q_l \Big[\sum_{n=0}^{\infty} \pi_n^l \sum_{i\in[K]} \lambda_l P_i \sigma^l(S_i|n)\Big] \\
&\overset{(b)}{=} \sum_{l=1}^{L} q_l \pi_n^l \Big[\sum_{n=0}^{\infty} \sum_{i\in[K]} \lambda_l P_i \sigma(S_i|n, \lambda_l)\Big] = \sum_{i\in[K]} \mathbb{E}_{\boldsymbol{\sigma}}[\Lambda P_i \xi(S_i)],
\end{aligned}$$

where the (a) follows because $\nu_j(S_i) = 1$ for all $j \geq i$; and (b) uses $\sigma(S_i|n, \lambda_l) = \sigma^l(S_i|n)$.

Now, to ensure that the customer equilibrium is given by $\boldsymbol{\nu}(S_i) = S_i$ for all $i \in [K+1]$, for any given $i \in [K+1]$, we need that $\mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_i) h_{i^+}(W)] \geq 0$ for all $i^+ \geq i$, and $\mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_i) h_{i^-}(W)] \leq 0$ for all $i^- < i$. By using the arguments in Eq. (18), the condition $\mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_i) h_{i^+}(W)] \geq 0$ for all $i^+ \geq i$ is equivalent to $\mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_i) h_i(W)] \geq 0$. And similarly, $\mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_i) h_{i^-}(W)] \leq 0$ for all $i^- < i$ is equivalent to $\mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_i) h_{i-1}(W)] \leq 0$. This gives us the IC constraints in the Throughput C-MDP (4). □

Appendix D: Proof of results in Section 4 and Section 5

D.1. Proof of Proposition 3

To prove the result in Proposition 3, we use the LP formulation presented in Corollary 2. Here, we use similar notations as presented in Section 6.1. Under no demand variation, we drop the script l, and we define the variable $\boldsymbol{\phi}$ as follows.

$$\phi_{n,i} := \mathbb{P}_{\boldsymbol{\sigma}}(W = n, \xi(S_i) = 1) = \pi_n \sigma(S_i|n).$$

Further, for all $i \in [K+1]$, $\boldsymbol{\phi}_i = \{\phi_{n,i}\}_{n\geq 0}$ and $\boldsymbol{\phi} = \{\boldsymbol{\phi}_i\}_{i\in[K+1]}$. Thus, by using Corollary 2, we have the following linear program.

$$\begin{aligned}
\max_{\boldsymbol{\phi}} \quad & \sum_{i=1}^{K} \lambda P_i \langle \boldsymbol{\phi}_i, \mathbb{1} \rangle \\
\text{subject to} \quad & \langle \boldsymbol{\phi}_i, \mathbf{h}_i \rangle \geq 0,\ \forall i \in [K], \ \text{ and } \ \langle \boldsymbol{\phi}_{i+1}, \mathbf{h}_i \rangle \leq 0,\ \forall i \in [K], \\
& \sum_{i=1}^{K+1} \phi_{n+1,i} - \sum_{i=1}^{K} \lambda P_i \phi_{n,i} = 0,\ \forall n \geq 0, \ \text{ and } \ \sum_{i=1}^{K+1} \langle \boldsymbol{\phi}_i, \mathbb{1} \rangle = 1, \\
& \phi_{n,i} \geq 0,\ \forall n \geq 0 \text{ and } \forall i \in [K+1].
\end{aligned} \tag{20}$$

For simplicity of notation, we consider

$$C_i := \frac{c_i}{b_i},\ \forall i \in [K], \text{ and } \tilde{\mathbf{h}}_i = \frac{1}{b_i}\mathbf{h}_i,$$

that is, $\tilde{h}_i(n) = C_i - n$. Note that, under these notations, C_i's are increasing in i, and $\tilde{h}_i(\cdot)$ are *ordered*, i.e., $\tilde{h}_1(n) < \tilde{h}_2(n) < \cdots < \tilde{h}_K(n)$ for all $n \geq 0$. Further, for any $i \in [K]$, the constraints $\langle \boldsymbol{\phi}_i, \mathbf{h}_i \rangle \geq 0$ and $\langle \boldsymbol{\phi}_{i+1}, \mathbf{h}_i \rangle \leq 0$ are equivalent to $\langle \boldsymbol{\phi}_i, \tilde{\mathbf{h}}_i \rangle \geq 0$ and $\langle \boldsymbol{\phi}_{i+1}, \tilde{\mathbf{h}}_i \rangle \leq 0$, respectively.

Next, we argue that the constraints given by $\langle \boldsymbol{\phi}_{i+1}, \tilde{\mathbf{h}}_i \rangle \leq 0,\ \forall i \in [K]$ are non-binding. To prove this, we use a contradiction based argument. Let $\boldsymbol{\phi}^*$ be an optimal solution of the LP in (20) and suppose

$$i_1 := \min\left\{ i \in [K] : \boldsymbol{\phi}^*_{i+1} \neq \mathbf{0}, \langle \boldsymbol{\phi}^*_{i+1}, \tilde{\mathbf{h}}_i \rangle = 0 \right\}.$$

Note that $C_i < C_{i+1}$, and so $\tilde{h}_i(n) < \tilde{h}_{i+1}(n)$ for all $n \geq 0$ and $i \in [K-1]$. Suppose $N_{i_1} := \max\{n : \tilde{h}_{i_1}(n) \geq 0\} = \lfloor C_{i_1} \rfloor$. Then, there exists an $n_1 \leq N_{i_1}$ such that $\phi^*_{n_1, i_1+1} > 0$, as otherwise $\langle \boldsymbol{\phi}^*_{i_1+1}, \tilde{\mathbf{h}}_{i_1} \rangle = \sum_{n > N_{i_1}} \phi^*_{n, i_1+1} \tilde{h}_{i_1}(n) < 0$, which violates the fact that $\langle \boldsymbol{\phi}^*_{i_1+1}, \tilde{\mathbf{h}}_{i_1} \rangle = 0$. Further, suppose $i_2 = \min\{i \in [K] : n_1 \leq C_i\}$. Note that $i_2 \leq i_1$ (as $n_1 \leq N_{i_1} \leq C_{i_1}$) and also $C_{i_2 - 1} < n_1 \leq C_{i_2}$ (as i_2 is the smallest i such that $n_1 \leq C_i$), where we set $C_0 = -\infty$.

From the above construction, we conclude that

$$C_{i_2-1} < n_1 \leq C_{i_2} \leq C_{i_1} < C_{i_1+1}, \text{ and, } \tilde{h}_{i_2-1}(n_1) < 0 \leq \tilde{h}_{i_2}(n_1) \leq \tilde{h}_{i_1}(n_1) < \tilde{h}_{i_1+1}(n_1). \tag{21}$$

Now, we construct another solution $\tilde{\boldsymbol{\phi}}$ as follows.

$$\tilde{\phi}_{n_1, i_1+1} = \frac{1}{Z}(1-\beta)\phi^*_{n_1, i_1+1}, \text{ and } \tilde{\phi}_{n_1, i_2} = \frac{1}{Z}\left[\phi^*_{n_1, i_2} + \beta \phi^*_{n_1, i_1+1}\right],$$
$$\tilde{\phi}_{n_1+1, K+1} = \frac{1}{Z}\left[\phi^*_{n_1+1, K+1} + \lambda\beta\phi^*_{n_1, i_1+1}(P_{i_2} - P_{i_1+1})\right],$$

and for all other values n and i, we have $\tilde{\phi}_{n,i} = \frac{1}{Z}\phi^*_{n,i}$. We claim that there exists $\beta > 0$ and $Z \geq 1$ such that $\tilde{\boldsymbol{\phi}}$ forms a feasible solution.

Claim: $\tilde{\boldsymbol{\phi}}$ satisfies the constraints corresponding to the flow balance equations.

Note that

$$\sum_{i=1}^{K+1} \langle \tilde{\boldsymbol{\phi}}_i, \mathbb{1} \rangle = \frac{1}{Z}\Big[\sum_{i=1}^{K+1} \langle \boldsymbol{\phi}^*_i, \mathbb{1} \rangle + \lambda\beta\phi^*_{n_1, i_1+1}(P_{i_2} - P_{i_1+1})\Big]$$
$$= \frac{1}{Z}\Big[1 + \lambda\beta\phi^*_{n_1, i_1+1}(P_{i_2} - P_{i_1+1})\Big],$$

and so by choosing $Z = 1 + \lambda\beta\phi^*_{n_1, i_1+1}(P_{i_2} - P_{i_1+1}) > 1$ (as $P_{i_2} > P_{i_1+1}$ using $i_2 < i_1 + 1$), we have that $\tilde{\boldsymbol{\phi}}$ satisfies $\sum_{i=1}^{K+1} \langle \tilde{\boldsymbol{\phi}}_i, \mathbb{1} \rangle = 1$ for all $\beta \in (0,1)$. Now, simply by the construction of $\tilde{\boldsymbol{\phi}}$, for all $n \geq 0$,

$$\sum_{i=1}^{K} \tilde{\phi}_{n+1,i} - \sum_{i=1}^{K} \lambda P_i \tilde{\phi}_{n,i} = \frac{1}{Z}\Big[\sum_{i=1}^{K+1} \phi^*_{n+1,i} - \sum_{i=1}^{K} \lambda P_i \phi^*_{n,i}\Big] = 0.$$

Finally, for all $\beta \in [0,1]$, $\tilde{\phi}_{n,i} \geq 0$ for all $n \geq 0$ and $i \in [K+1]$.

From the above arguments, $\tilde{\boldsymbol{\phi}}$ satisfies the constraints corresponding to the flow balance equations, that is, the last three constraints in LP (20).

Claim: $\tilde{\boldsymbol{\phi}}$ satisfies the IC constraints for some $\beta \in (0,1)$.

- Consider the first constraint in the LP in (20). We have,

$$\langle \tilde{\boldsymbol{\phi}}_i, \tilde{\mathbf{h}}_i \rangle = \frac{1}{Z} \begin{cases} \langle \boldsymbol{\phi}^*_{i_1}, \tilde{\mathbf{h}}_{i_2} \rangle + \beta \phi^*_{n_1, i_1+1} \tilde{h}_{i_2}(n_1) & \text{for } i = i_2 \\ \langle \boldsymbol{\phi}^*_{i_1+1}, \tilde{\mathbf{h}}_{i_1+1} \rangle - \beta \phi^*_{n_1, i_1+1} \tilde{h}_{i_1+1}(n_1) & \text{for } i = i_1+1, \text{ if } i_1 < K \\ \langle \boldsymbol{\phi}^*_i, \tilde{\mathbf{h}}_i \rangle & \text{otherwise} \end{cases}$$
$$\geq \frac{1}{Z} \begin{cases} \langle \boldsymbol{\phi}^*_{i_1+1}, \tilde{\mathbf{h}}_{i_1+1} \rangle - \beta \phi^*_{n_1, i_1+1} \tilde{h}_{i_1+1}(n_1) & \text{for } i = i_1+1, \text{ if } i_1 < K \\ 0 & \text{otherwise,} \end{cases}$$

where we use $\tilde{h}_{i_2}(n_1) \geq 0$ by construction of n_1 (see Eq. (21)). From the above argument, for $i_1 = K$, we trivially have $\langle \tilde{\boldsymbol{\phi}}_i, \tilde{\mathbf{h}}_i \rangle \geq 0$ for all $i \in [K]$. Now, for $i_1 < K$,

$$\langle \boldsymbol{\phi}^*_{i_1+1}, \tilde{\mathbf{h}}_{i_1+1} \rangle = \langle \boldsymbol{\phi}^*_{i_1+1}, \tilde{\mathbf{h}}_{i_1} \rangle + \langle \boldsymbol{\phi}^*_{i_1+1}, \tilde{\mathbf{h}}_{i_1+1} - \tilde{\mathbf{h}}_{i_1} \rangle = \langle \boldsymbol{\phi}^*_{i_1+1}, \tilde{\mathbf{h}}_{i_1+1} - \tilde{\mathbf{h}}_{i_1} \rangle > 0, \tag{22}$$

where the second equality holds as $\langle \boldsymbol{\phi}^*_{i_1+1}, \tilde{\mathbf{h}}_{i_1} \rangle = 0$ by definition of i_1. We also have $\tilde{h}_{i_1+1}(n_1) > 0$ by construction of n_1 (see Eq. (21)). So, we can pick

$$0 < \beta \leq \min \left\{ 1, \frac{\langle \boldsymbol{\phi}^*_{i_1+1}, \tilde{\mathbf{h}}_{i_1+1} \rangle}{\phi^*_{n_1, i_1+1} \tilde{h}_{i_1+1}(n_1)} \right\},$$

which gives us that $\langle \tilde{\boldsymbol{\phi}}_i, \tilde{\mathbf{h}}_i \rangle \geq 0$ for all $i \in [K]$.

- For $i \in [K]$, we have

$$\langle \tilde{\boldsymbol{\phi}}_{i+1}, \tilde{\mathbf{h}}_i \rangle = \frac{1}{Z} \begin{cases} \langle \boldsymbol{\phi}^*_{i_2}, \tilde{\mathbf{h}}_{i_2-1} \rangle + \beta \phi^*_{n_1, i_1+1} \tilde{h}_{i_2-1}(n_1) & \text{for } i = i_2 - 1, \text{ if } i_2 > 1 \\ \langle \boldsymbol{\phi}^*_{i_1+1}, \tilde{\mathbf{h}}_{i_1} \rangle - \beta \phi^*_{n_1, i_1+1} \tilde{h}_{i_1}(n_1) & \text{for } i = i_1, \text{ if } i_1 > 1 \\ \langle \boldsymbol{\phi}^*_{i+1}, \tilde{\mathbf{h}}_i \rangle & \text{otherwise} \end{cases}$$
$$\overset{(a)}{\leq} \frac{1}{Z} \begin{cases} \beta \phi^*_{n_1, i_1+1} \tilde{h}_{i_2-1}(n_1) & \text{for } i = i_2 - 1, \text{ if } i_2 > 1 \\ -\beta \phi^*_{n_1, i_1+1} \tilde{h}_{i_1}(n_1) & \text{for } i = i_1, \text{ if } i_1 > 1 \\ 0 & \text{otherwise} \end{cases}$$
$$\overset{(b)}{\leq} 0$$

where (a) follows by using $\langle \boldsymbol{\phi}^*_{i+1}, \tilde{\mathbf{h}}_i \rangle \leq 0, \forall i \in [K]$; and (b) follows by using $\tilde{h}_{i_1}(n_1) \geq 0 > \tilde{h}_{i_2-1}(n_1)$ (see Eq. (21)). This shows that there exists $\beta \in (0,1)$ such that $\tilde{\boldsymbol{\phi}}$ satisfies both the IC constraints.

The above arguments imply that $\tilde{\boldsymbol{\phi}}$ is a feasible solution for β small enough. We also note that the objective value of the $\tilde{\boldsymbol{\phi}}$ is higher than that of $\boldsymbol{\phi}^*$ for any $\beta > 0$ as

$$\sum_{i=1}^{K} \lambda P_i \langle \tilde{\boldsymbol{\phi}}_i, \mathbb{1} \rangle - \sum_{i=1}^{K} \lambda P_i \langle \boldsymbol{\phi}^*_i, \mathbb{1} \rangle = \sum_{n=0}^{\infty} \sum_{i=1}^{K} \lambda P_i \tilde{\phi}_{n,i} - \sum_{n=0}^{\infty} \sum_{i=1}^{K} \lambda P_i \phi^*_{n,i}$$

$$
\begin{aligned}
&= \sum_{n=0}^{\infty}\sum_{i=1}^{K+1}\tilde{\phi}_{n+1,i} - \sum_{n=0}^{\infty}\sum_{i=1}^{K+1}\phi^*_{n+1,i} \\
&= \Big[1 - \sum_{i=1}^{K+1}\tilde{\boldsymbol{\phi}}_{0,i}\Big] - \Big[1 - \sum_{i=1}^{K+1}\boldsymbol{\phi}^*_{0,i}\Big] \\
&= \Big(1 - \frac{1}{Z}\Big)\sum_{i=1}^{K+1}\boldsymbol{\phi}^*_{0,i} \\
&> 0,
\end{aligned}
$$

where the last inequality follows as $Z > 1$. As such, we have a new solution $\tilde{\boldsymbol{\phi}}$ which is feasible and has a higher objective value than $\boldsymbol{\phi}^*$. This leads to a contradiction, as $\boldsymbol{\phi}^*$ is the optimal solution. This implies that there does not exist i_1 such that $\boldsymbol{\phi}^*_{i_1+1} \neq \mathbf{0}$, and $\langle \boldsymbol{\phi}^*_{i_1+1}, \tilde{\mathbf{h}}_{i_1}\rangle = 0$. and so, $\langle \boldsymbol{\phi}^*_{i+1}, \tilde{\mathbf{h}}_i\rangle < 0$ for all $i \in [K]$. This completes the proof. □

D.2. Proof of Proposition 5

In case of homogeneous customers $(K = 1)$ where customers are unaware of the realized arrival rate, the Throughput LP (6) reduces to the following.

$$
\begin{aligned}
\max_{\boldsymbol{\phi}} \quad & \sum_{l=1}^{L} q_l \lambda_l \langle \boldsymbol{\phi}^l_1, \mathbb{1}\rangle \\
\text{s.t.} \quad & \sum_{l=1}^{L} q_l \langle \boldsymbol{\phi}^l_1, \mathbf{h}\rangle \geq 0, \text{ and } \sum_{l=1}^{L} q_l \langle \boldsymbol{\phi}^l_2, \mathbf{h}\rangle \leq 0, \\
& \phi^l_{n+1,1} + \phi^l_{n+1,2} = \lambda_l \phi^l_{n,1}, \ \forall n \geq 0, \text{ and } \forall l \in [L], \\
& \langle \boldsymbol{\phi}^l_1, \mathbb{1}\rangle + \langle \boldsymbol{\phi}^l_2, \mathbb{1}\rangle = 1, \ \forall l \in [L], \\
& \phi^l_{n,1}, \phi^l_{n,2} \geq 0, \ \forall n \geq 0 \text{ and } \forall l \in [L],
\end{aligned}
\tag{23}
$$

where $\mathbf{h} = \{c - bn\}_{\{n \geq 0\}}$. In contrast, in the scenario where the customer know the realized arrival rate (say $\Lambda = \lambda_l$), the server uses the optimal policy corresponding to case when arrival rate realization is fixed and given by λ_l. As such, the server solves the following LP, which is the reduction of Throughput LP (6) when customer are homogeneous and the arrival rate is λ_l.

$$
\begin{aligned}
\max_{\boldsymbol{\phi}} \quad & \lambda_l \langle \boldsymbol{\phi}^l_1, \mathbb{1}\rangle \\
\text{s.t.} \quad & \langle \boldsymbol{\phi}^l_1, \mathbf{h}\rangle \geq 0, \text{ and } \langle \boldsymbol{\phi}^l_2, \mathbf{h}\rangle \leq 0, \\
& \phi^l_{n+1,1} + \phi^l_{n+1,2} = \lambda_l \phi^l_{n,1}, \ \forall n \geq 0, \text{ and } \langle \boldsymbol{\phi}^l_1, \mathbb{1}\rangle + \langle \boldsymbol{\phi}^l_2, \mathbb{1}\rangle = 1, \ \forall l \in [L], \\
& \phi^l_{n,1}, \phi^l_{n,2} \geq 0, \ \forall n \geq 0,
\end{aligned}
\tag{24}
$$

where, with abuse of notation, we use the same variables as in LP in (23). Suppose $\{\tilde{\boldsymbol{\phi}}^l = \{\tilde{\boldsymbol{\phi}}^l_1, \tilde{\boldsymbol{\phi}}^l_2\}\}_{\{l \in [L]\}}$ is set of the optimal solutions of LP in (24) for the set of arrival rates $\{\lambda_1, \lambda_2, \ldots, \lambda_L\}$. Then, $\langle \tilde{\boldsymbol{\phi}}^l_1, \mathbf{h}\rangle \geq 0, \ \forall l \in$

$[L]$, which implies that $\sum_{l=1}^{L} q_l \langle \tilde{\boldsymbol{\phi}}_1^l, \mathbf{h}\rangle \geq 0$. Similarly, $\langle \tilde{\boldsymbol{\phi}}_2^l, \mathbf{h}\rangle \leq 0,\ \forall l \in [L]$ implies that $\sum_{l=1}^{L} q_l \langle \tilde{\boldsymbol{\phi}}_2^l, \mathbf{h}\rangle \leq 0$. Note that all other constraints in LP in (23) match the constraints in LP in (24), and so they also remain satisfied by $\left\{\tilde{\boldsymbol{\phi}}^l = \{\tilde{\boldsymbol{\phi}}_1^l, \tilde{\boldsymbol{\phi}}_2^l\}\right\}_{\{l\in[L]\}}$. As such, $\left\{\tilde{\boldsymbol{\phi}}^l = \{\tilde{\boldsymbol{\phi}}_1^l, \tilde{\boldsymbol{\phi}}_2^l\}\right\}_{\{l\in[L]\}}$ is a feasible solution of the LP in (23). This implies that the optimal objective value of the LP in (23) (averaged over the distribution of arrival rates) is at least as large as the optimal objective value of the LP in (24).

D.3. Proof of Proposition 4

From the result in Theorem 1, we get that the optimal policy is ordered laminar. Next, we provide the additional arguments required to show that the policy is monotone. Note that, under homogeneous customer types, we only have two signals S_1 and S_2. For simplicity of the notation, we consider a fixed $l \in [L]$, and we drop the superscript and the subscript corresponding to the arrival rate. Let n_1 be the smallest n such that $\sigma^l(S_2|n) > 0$.

Note that at $n = n_1 + 1$, if $\sigma^l(S_2|n) = 1$, then without loss of generality, we can consider an equivalent policy where $\sigma^l(S_2|n) = 1$ for all $n \geq n_1 + 1$, which gives us that $\sigma^l(S_1|n) = 0$ for all $n \geq n_1 + 1$ and so the policy have a threshold structure by choosing $T_l = n_1 + \sigma^l(S_2|n_1)$.

Next, we consider the case where $\sigma^l(S_2|n) < 1$ for $n = n_1 + 1$ (implying $\sigma^l(S_1|n) > 0$ at $n = n_1 + 1$). Then, as the optimal policy is laminar, we get that $\sigma^l(S_2|n) = 0$ for all $n > n_1 + 1$, as otherwise $\sigma^l(S_1|n) = 0$ at $n = n_1 + 1$ using the statement in Lemma 1. Before presenting the next step, we recall that, for any policy $\boldsymbol{\sigma}$ (with steady-state distribution tuple $\boldsymbol{\Pi}$), the variable $\boldsymbol{\phi}$ denotes $\phi_{n,i}^l = \pi_n^l \sigma^l(S_i|n),\ \text{ for all } n \geq 0, l \in [L],\ \text{and } i \in \{1,2\}$.

Next, we construct a new threshold policy $\tilde{\boldsymbol{\sigma}}^l$ with threshold T_l such that $\mathbb{E}_{\tilde{\boldsymbol{\sigma}}}[\xi(S_1)|\Lambda = \lambda_l]$ remains the same (where $\tilde{\boldsymbol{\sigma}}^{l'} = \boldsymbol{\sigma}^{l'}$ for $l' \neq l$). Mathematically, T_l is the solution of

$$\tilde{\pi}_0^l \Big[\sum_{n=0}^{\lfloor T_l \rfloor} \lambda_l^n + \{T_l\}\lambda_l^{\lfloor T_l \rfloor + 1}\Big] = 1, \ \text{ and } \ \tilde{\pi}_0^l \Big[\sum_{n=0}^{\lfloor T_l \rfloor - 1} \lambda_l^n + \{T_l\}\lambda_l^{\lfloor T_l \rfloor}\Big] = \mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_1)|\Lambda = \lambda_l], \tag{25}$$

where the first equation come from the constraint $\sum_{n\geq 0} \tilde{\pi}_n^l = 1$ and the second equation is required to ensure $\mathbb{E}_{\tilde{\boldsymbol{\sigma}}}[\xi(S_1)|\Lambda = \lambda_l] = \mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_1)|\Lambda = \lambda_l]$, where we use the fact that $\tilde{\boldsymbol{\sigma}}^l$ is a threshold policy with threshold T_l and so,

$$\tilde{\sigma}^l(S_1|n) = \mathbb{1}_{\{n < \lfloor T_l \rfloor\}} + \{T_l\}\mathbb{1}_{\{n = \lfloor T_l \rfloor\}}.$$

Note that a T_l that solves the Eq. (25) exists because $\mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_1)|\Lambda = \lambda_l] < 1$. Now, we claim that the IC constraints remain satisfied under the new policy. By construction of $\tilde{\boldsymbol{\sigma}}^l$, we have $\mathbb{E}_{\tilde{\boldsymbol{\sigma}}}[\xi(S_1)|\Lambda = \lambda_l] = \mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_1)|\Lambda = \lambda_l]$ (and same for S_2). However, by using the threshold policy, the signal S_1 is now provided at smaller queue length values and so, the wait time under the signal S_1 is now reduced, meaning $\mathbb{E}_{\tilde{\boldsymbol{\sigma}}}[W\xi(S_1)|\Lambda = \lambda_l] < \mathbb{E}_{\boldsymbol{\sigma}}[W\xi(S_1)|\Lambda = \lambda_l]$. As such, the first IC constraint remains satisfied. Further, under the policy $\boldsymbol{\sigma}^l$, the signal S_2 was given only at the state n_1 and possibly at $n_1 + 1$. By the construction of

$\tilde{\boldsymbol{\sigma}}^l$, the probability of giving the signal S_2 at n_1 is reduced while the overall probability of giving signal S_2 remains the same. This implies that the expected wait time under the signal S_2 is now increased, meaning $\mathbb{E}_{\tilde{\boldsymbol{\sigma}}}[W\xi(S_2)|\Lambda=\lambda_l] > \mathbb{E}_{\boldsymbol{\sigma}}[W\xi(S_2)|\Lambda=\lambda_l]$. As such, the second IC constraint also remains satisfied.

This gives a threshold policy under which the system throughput remains the same while the IC constraints are still satisfied. And so, it suffices for the server to choose a threshold policy.

Appendix E: Proof of results in Section 6

E.1. Proof of Lemma 2

We say that a policy $\boldsymbol{\sigma}$ is feasible if the system is stable, i.e., the stationary distribution $\boldsymbol{\pi}^l$ exists for all $l \in [L]$, and the constraints in the throughput optimization problem (4) are satisfied. In this case, we define the variables $(\mathbf{p}, \mathbf{w}) = \{p^l(S_i), w^l(S_i)\}_{i\in[K+1], l\in[L]}$ as follows.

$$p^l(S_i) := \sum_{n\geq 0} \pi_n^l \sigma^l(S_i|n), \qquad\qquad w^l(S_i) := \sum_{n\geq 0} n\pi_n^l \sigma^l(S_i|n).$$

Suppose $\mathcal{X}$ be the set of $(\mathbf{p}, \mathbf{w})$ such that there exists a feasible policy $\boldsymbol{\sigma}$ satisfying $(\mathbf{p}, \mathbf{w}) = \chi(\boldsymbol{\sigma})$. Next, we prove that the set $\mathcal{X}$ is compact and convex.

<u>The set $\mathcal{X}$ is bounded</u>

In order to show that the set $\mathcal{X}$ is bounded, we provide an upper and a lower bound on the variables $p^l(S_i)$ and $w^l(S_i)$ for all $i \in [K+1]$ and $l \in [L]$.

Consider a given $l \in [L]$. As $\sigma^l(S_i|n) \in [0,1]$ and $\sum_{n\geq 0} \pi_n^l = 1$, we easily get that $p^l(S_i) \in [0,1]$ for all $i \in [K+1]$. Next, we use the constraints in the throughput optimization problem (4) to create a bound on the variable $\mathbf{w}$ as follows. The first IC constraint in (4) gives us that, for all $i \in [K]$, $\sum_{l\in[L]} \sum_{n=0}^{\infty} q_l \pi_n^l \sigma^l \big(c_i - b_i n\big) \geq 0$. Now, by using the definition of $p^l(S_i)$ and $w^l(S_i)$, we can rewrite the above IC constraints as $\sum_{l\in[L]} q_l \big(c_i p^l(S_i) - b_i w^l(S_i)\big) \geq 0$. This gives us that, for all $i \in [K]$,

$$\sum_{l\in[L]} q_l w^l(S_i) \leq \frac{c_i}{b_i} \sum_{l\in[L]} q_l p^l(S_i) \leq \frac{c_i}{b_i},$$

where we use the fact that $p^l(S_i) \leq 1$. As such, $0 \leq w^l(S_i) \leq \frac{c_i}{q_l b_i}$ for all $i \in [K]$ and $l \in [L]$. Note that the above argument only provides a bound on $w^l(S_i)$ for $i \in [K]$. For $i = K+1$,

$$w^l(S_{K+1}) = \sum_{n\geq 0} n\pi_n^l \sigma^l(S_{K+1}|n) \leq \sum_{n\geq 1} n\pi_n^l \overset{(a)}{=} \lambda_l \sum_{i=1}^{K} \sum_{n\geq 1} n\pi_{n-1}^l P_i \sigma^l(S_i|n-1),$$

where (a) follows using the flow balance equation, i.e. $\pi_n^l = \lambda_l \pi_{n-1}^l \sum_{i=1}^{K} P_i \sigma^l(S_i|n-1)$ for all $n \geq 1$. This gives us that

$$\begin{aligned} w^l(S_{K+1}) \leq \lambda_l \sum_{i=1}^{K} \sum_{n\geq 0} (n+1)\pi_n^l P_i \sigma^l(S_i|n) &\overset{(a)}{=} \lambda_l \sum_{i=1}^{K} P_i (w^l(S_i) + p^l(S_i)) \\ &\overset{(b)}{\leq} \lambda_l \sum_{i=1}^{K} P_i \Big(1 + \frac{c_i}{q_l p_i}\Big), \end{aligned}$$

where (a) follows by definition of $p^l(S_i)$ and $w^l(S_i)$; and (b) follows by using $p^l(S_i) \leq 1$ and $w^l(S_i) \leq \frac{c_i}{q_l p_i}$ for all $i \in [K]$. As this holds for all $l \in [L]$, we get that $(\mathbf{p}, \mathbf{w})$ lies in a bounded set, i.e., the set $\mathcal{X}$ is bounded.

The set $\mathcal{X}$ is convex:

Suppose $(\mathbf{p}, \mathbf{w})$ and $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$ are two elements in the set $\mathcal{X}$. Let $\gamma \in [0,1]$ and $(\tilde{\mathbf{p}}, \tilde{\mathbf{w}}) = \gamma(\mathbf{p}, \mathbf{w}) + (1-\gamma)(\hat{\mathbf{p}}, \hat{\mathbf{w}})$. That is, $\tilde{p}^l(S_i) = \gamma p^l(S_i) + (1-\gamma)\hat{p}^l(S_i)$ (and same for $\tilde{w}^l(S_i)$) for all $i \in [K+1]$ and $l \in [L]$. We need to show that $(\tilde{\mathbf{p}}, \tilde{\mathbf{w}})$ lies in the set $\mathcal{X}$.

As $(\mathbf{p}, \mathbf{w}) \in \mathcal{X}$, there exists a feasible policy $\boldsymbol{\sigma}$, such that $(\mathbf{p}, \mathbf{w}) = \chi(\boldsymbol{\sigma})$. Similarly, there exists a feasible policy $\hat{\boldsymbol{\sigma}}$, such that $(\hat{\mathbf{p}}, \hat{\mathbf{w}}) = \chi(\hat{\boldsymbol{\sigma}})$. Suppose $\boldsymbol{\Pi} = \{\boldsymbol{\pi}^l\}_{l \in [L]}$ and $\hat{\boldsymbol{\Pi}} = \{\hat{\boldsymbol{\pi}}^l\}_{l \in [L]}$ are the tuples of steady state distribution under the policy $\boldsymbol{\sigma}$ and $\hat{\boldsymbol{\sigma}}$ respectively.

Next, we define a new policy $\tilde{\boldsymbol{\sigma}}$ as follows. For all $i \in [K+1]$, $l \in [L]$ and $n \geq 0$, if $\gamma\pi_n^l + (1-\gamma)\hat{\pi}_n^l > 0$, we choose

$$\tilde{\sigma}^l(S_i|n) = \frac{\gamma\pi_n^l \sigma^l(S_i|n) + (1-\gamma)\hat{\pi}_n^l \hat{\sigma}^l(S_i|n)}{\gamma\pi_n^l + (1-\gamma)\hat{\pi}_n^l},$$

and when $\gamma\pi_n^l + (1-\gamma)\hat{\pi}_n^l = 0$, we choose $\tilde{\sigma}^l(S_i|n) = 0$ for $i \in [K]$ and $\tilde{\sigma}^l(S_{K+1}|n) = 1$. Suppose $\tilde{\pi}_n^l := \gamma\pi_n^l + (1-\gamma)\hat{\pi}_n^l$ for all $n \geq 0$ and $l \in [L]$. Then, we have

$$\tilde{\pi}_n^l \tilde{\sigma}^l(S_i|n) = \gamma\pi_n^l \sigma^l(S_i|n) + (1-\gamma)\hat{\pi}_n^l \hat{\sigma}^l(S_i|n), \; \forall i \in [K+1], l \in [L], \text{ and } n \geq 0, \tag{26}$$

where the equality holds even if $\tilde{\pi}_n^l = 0$. Recall that, for all $l \in [L]$, $\boldsymbol{\pi}^l$ and $\hat{\boldsymbol{\pi}}^l$ satisfy the respective flow balance equations. As such, for all $l \in [L]$,

$$\sum_{n \geq 0} \tilde{\pi}_n^l = \gamma \sum_{n \geq 0} \pi_n^l + (1-\gamma) \sum_{n \geq 0} \hat{\pi}_n^l = 1,$$

and, by using Eq. (26), for all $l \in [L]$ and $n \geq 0$,

$$\begin{aligned}
\lambda_l \sum_{i=1}^{K} P_i \tilde{\pi}_n^l \tilde{\sigma}^l(S_i|n) &= \gamma\lambda_l \sum_{i=1}^{K} P_i \pi_n^l \sigma^l(S_i|n) + (1-\gamma)\lambda_l \sum_{i=1}^{K} P_i \hat{\pi}_n^l \hat{\sigma}^l(S_i|n) \\
&= \gamma\pi_{n+1}^l + (1-\gamma)\hat{\pi}_{n+1}^l \\
&= \tilde{\pi}_{n+1}^l.
\end{aligned}$$

This shows that $\tilde{\boldsymbol{\Pi}} = \{\tilde{\boldsymbol{\pi}}^l\}_{l \in [L]}$, where $\tilde{\boldsymbol{\pi}}^l = \{\tilde{\pi}_n^l\}_{n \geq 0}$, is the steady-state distribution tuple under the policy $\tilde{\boldsymbol{\sigma}}$.

Next, we argue that the policy $\tilde{\boldsymbol{\sigma}}$ satisfy the IC constraints. As $\tilde{\boldsymbol{\Pi}} = \{\tilde{\boldsymbol{\pi}}^l\}_{l \in [L]}$ is the steady-state distribution tuple under $\tilde{\boldsymbol{\sigma}}$, for all $i \in [K]$,

$$\mathbb{E}_{\tilde{\boldsymbol{\sigma}}}[\xi(S_i) h_i(W)] = \sum_{l \in [L]} \sum_{n=0}^{\infty} q_l \tilde{\pi}_n^l \tilde{\sigma}^l(S_i|n)\big(c_i - b_i n\big)$$

$$\stackrel{(a)}{=} \gamma \sum_{l \in [L]} \sum_{n=0}^{\infty} q_l \pi_n^l \sigma^l(S_i|n)\big(c_i - b_i n\big)$$
$$+ (1-\gamma) \sum_{l \in [L]} \sum_{n=0}^{\infty} q_l \hat{\pi}_n^l \hat{\sigma}^l(S_i|n)\big(c_i - b_i n\big)$$
$$= \gamma \mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_i)h_i(W)] + (1-\gamma)\mathbb{E}_{\hat{\boldsymbol{\sigma}}}[\xi(S_i)h_i(W)] \stackrel{(b)}{\geq} 0,$$

where (a) uses Eq. (26); and (b) follows as both $\boldsymbol{\sigma}$ and $\hat{\boldsymbol{\sigma}}$ satisfy the IC constraints. Similarly, for all $i \in [K]$, $\mathbb{E}_{\tilde{\boldsymbol{\sigma}}}[\xi(S_{i+1})h_i(W)] \leq 0$, and so $\tilde{\boldsymbol{\sigma}}$ is feasible policy.

Finally, we claim that $(\tilde{\mathbf{p}}, \tilde{\mathbf{w}}) = \chi(\tilde{\boldsymbol{\sigma}})$. For all $i \in [K+1]$ and $l \in [L]$,

$$\tilde{p}^l(S_i) = \gamma p^l(S_i) + (1-\gamma)\hat{p}^l(S_i)$$
$$= \gamma \sum_{n=0}^{\infty} \pi_n^l \sigma^l(S_i|n) + (1-\gamma)\sum_{n=0}^{\infty} \hat{\pi}_n^l \hat{\sigma}^l(S_i|n) = \sum_{n=0}^{\infty} \tilde{\pi}_n^l \tilde{\sigma}^l(S_i|n),$$

where the last equality follows by using Eq. (26). By using the same argument for $w^l(S_i)$'s, we get that $(\tilde{\mathbf{p}}, \tilde{\mathbf{w}}) = \chi(\tilde{\boldsymbol{\sigma}})$, where $\tilde{\boldsymbol{\sigma}}$ is a feasible policy. As such $(\tilde{\mathbf{p}}, \tilde{\mathbf{w}}) \in \mathcal{X}$, and so $\mathcal{X}$ is a convex set.

<u>The set $\mathcal{X}$ is closed:</u>

In order to prove this part, we use the notations corresponding to the LP formulation of the throughput optimization problem, as presented in Section 6.1. Recall that, for any policy $\boldsymbol{\sigma}$ (with steady-state distribution tuple $\boldsymbol{\Pi}$), we define the variable $\boldsymbol{\phi}$ as $\phi_{n,i}^l := \pi_n^l \sigma^l(S_i|n)$. Further, for all $l \in [L]$, $\boldsymbol{\phi}_i^l = \{\phi_{n,i}^l\}_{n\geq 0}$, $\boldsymbol{\phi}^l = \{\boldsymbol{\phi}_1^l, \boldsymbol{\phi}_2^l, \ldots, \boldsymbol{\phi}_{K+1}^l\}$, and $\boldsymbol{\phi} = \{\boldsymbol{\phi}^1, \boldsymbol{\phi}^2, \ldots, \boldsymbol{\phi}^L\}$. Note that, for any $l \in [L]$, $\boldsymbol{\phi}^l$ is a probability distribution over $\mathbb{Z}_+ \times [K+1]$. Using this notation, we have that $p^l(S_i) = \sum_{n\geq 0} \phi_{n,i}^l = \langle \boldsymbol{\phi}_i^l, \mathbb{1}\rangle$ and $w^l(S_i) = \sum_{n\geq 0} n\phi_{n,i}^l = \langle \boldsymbol{\phi}_i^l, \mathbf{I}\rangle$, where $\mathbf{I} = \{n\}_{n\geq 0}$. We use $\tilde{\chi}(\cdot)$ to denote the mapping $(\mathbf{p}, \mathbf{w}) = \tilde{\chi}(\boldsymbol{\phi})$. Note that the mapping $\tilde{\chi}(\cdot)$ is linear.

Suppose $\{(\mathbf{p}_{(m)}, \mathbf{w}_{(m)})\}_{m\geq 0}$ be a sequence of points in $\mathcal{X}$ such that the sequence converges to $(\mathbf{p}, \mathbf{w})$. For every $m \geq 0$, as $(\mathbf{p}_{(m)}, \mathbf{w}_{(m)}) \in \mathcal{X}$, there exists $\boldsymbol{\phi}_{(m)}$ such that $(\mathbf{p}_{(m)}, \mathbf{w}_{(m)}) = \tilde{\chi}(\boldsymbol{\phi}_{(m)})$.

Consider $l = 1$. As $\mathcal{X}$ is bounded, we know that $\sum_{i=1}^{K+1} w_{(m)}^1(S_i) = \sum_{i=1}^{K+1} \langle \boldsymbol{\phi}_{(m),i}^1, \mathbf{I}\rangle$ is bounded by a constant for all $m \geq 0$. This implies that the sequence of probability distribution $\{\boldsymbol{\phi}_{(m)}^1\}$ is tight. So, we can find a subsequence $\{m_k\}_{k\geq 0}$ such that $\{\boldsymbol{\phi}_{(m_k)}^1\}$ converges in distribution. Now, for $l = 2$, we can find a further subsequence of $\{m_k\}_{k\geq 0}$ along which $\{\boldsymbol{\phi}_{(m_k)}^2\}$ converges. By repeating this argument, we can construct a subsequence of $\{\boldsymbol{\phi}_{(m)}\}$ which converges to $\boldsymbol{\phi} = \{\boldsymbol{\phi}^1, \boldsymbol{\phi}^2, \ldots, \boldsymbol{\phi}^L\}$, where $\boldsymbol{\phi}^l$ is a probability distribution for all $l \in [L]$.

Next, we argue that $\boldsymbol{\phi}$ is feasible. With abuse of notation, suppose $\{\boldsymbol{\phi}_{(m)}\}$ converges to $\boldsymbol{\phi}$. As $\{\boldsymbol{\phi}_{(m)}^l\}$ is tight for all $l \in [L]$, we get that, for all $i \in [K+1]$,

$$\langle \boldsymbol{\phi}_i^l, \mathbb{1}\rangle = \lim_{m\to\infty} \langle \boldsymbol{\phi}_{(m),i}^l, \mathbb{1}\rangle = \lim_{m\to\infty} p_{(m)}^l(S_i) = p^l(S_i),$$

and similarly, $\langle \boldsymbol{\phi}_i^l, \mathbf{I}\rangle = w^l(S_i)$. This implies that $(\mathbf{p},\mathbf{w}) = \tilde{\chi}(\boldsymbol{\phi})$. Further, using similar argument as above, for all $i \in [K]$,

$$\langle \boldsymbol{\phi}_i^l, \mathbf{h}_i\rangle = \lim_{m\to\infty} \langle \boldsymbol{\phi}_{(m),i}^l, \mathbf{h}_i\rangle \geq 0,$$

where we use $\langle \boldsymbol{\phi}_{(m),i}^l, \mathbf{h}_i\rangle \geq 0$ as $\boldsymbol{\phi}_{(m)}$ is feasible. Similarly, for all $i \in [K]$, $\langle \boldsymbol{\phi}_{i+1}^l, \mathbf{h}_i\rangle \leq 0$. This shows that $\boldsymbol{\phi}$ also satisfies the IC constraints in the LP formulation in (6). Also, for all $l \in [L]$, $\sum_{i=1}^{K+1} \langle \boldsymbol{\phi}_i^l, \mathbb{1}\rangle = \lim_{m\to\infty} \sum_{i=1}^{K+1} \langle \boldsymbol{\phi}_{(m),i}^l, \mathbb{1}\rangle = 1$, and, for all $n \geq 0$, $l \in [L]$,

$$\sum_{i=1}^{K+1} \phi_{n+1,i}^l = \lim_{m\to\infty} \sum_{i=1}^{K+1} \phi_{(m),n+1,i}^l = \lim_{m\to\infty} \lambda_l \sum_{i=1}^{K+1} P_i \phi_{(m),n,i}^l = \lambda_l \sum_{i=1}^{K+1} P_i \phi_{n,i}^l.$$

As such, $\boldsymbol{\phi}$ also satisfies the flow balance equations in the LP formulation in (6). This implies $\boldsymbol{\phi}$ is feasible, implying that $(\mathbf{p},\mathbf{w}) = \tilde{\chi}(\boldsymbol{\phi}) \in \mathcal{X}$. This shows that $\mathcal{X}$ is closed.

<u>Slater's condition:</u>

Let $\{\mathcal{T}_i\}_{i\in[K+1]}$ be a monotone partition corresponding to the thresholds $x_i = \lfloor c_i/b_i \rfloor + 1$ for $i \in [K]$. Consider the monotone signaling mechanism where $\boldsymbol{\sigma}^l$ is monotone policy with respect to $\{\mathcal{T}_i\}_{i\in[K+1]}$ for all $l \in [L]$. Suppose $\boldsymbol{\Pi} = \{\boldsymbol{\pi}^l\}_{l\in[L]}$ denotes the corresponding steady-state distribution tuple.

Suppose $\mathcal{Z} = \{i \in [K] : \lfloor c_{i-1}/b_{i-1}\rfloor = \lfloor c_i/b_i\rfloor\}$. Note that for any $i \in \mathcal{Z}$, $\mathcal{T}_i = \emptyset$ and so $\sigma^l(S_i|n) = 0$ for all n. We argue that for the considered policy $\boldsymbol{\sigma} = \{\boldsymbol{\sigma}^l\}_{l\in[L]}$, for any $i \in [K]$, such that $i \notin \mathcal{Z}$ and $i > 1$, the IC constraints are strictly satisfied. We have

$$\begin{aligned}
&\sum_{n\geq 0} n\pi_n^l \sigma^l(S_i|n) \\
&\quad = \sum_{n=\lfloor c_{i-1}/b_{i-1}\rfloor+1}^{\lfloor c_i/b_i\rfloor} n\pi_n^l \sigma^l(S_i|n) \\
&\quad \in \left[\left(\left\lfloor \frac{c_{i-1}}{b_{i-1}}\right\rfloor + 1\right) \times \sum_{n=\lfloor c_{i-1}/b_{i-1}\rfloor+1}^{\lfloor c_i/b_i\rfloor} \pi_n^l \sigma^l(S_i|n), \left\lfloor \frac{c_i}{b_i}\right\rfloor \times \sum_{n=\lfloor c_{i-1}/b_{i-1}\rfloor+1}^{\lfloor c_i/b_i\rfloor} \pi_n^l \sigma^l(S_i|n)\right] \\
&\quad \subseteq \left(\frac{c_{i-1}}{b_{i-1}} \times \sum_{n\geq 0}^{\infty} \pi_n^l \sigma^l(S_i|n), \frac{c_i}{b_i} \times \sum_{n\geq 0}^{\infty} \pi_n^l \sigma^l(S_i|n)\right).
\end{aligned} \tag{27}$$

As such, for the policy $\boldsymbol{\sigma}$, $\mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_i)(c_i - b_iW)] > 0$ and $\mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_{i+1})(c_i - b_iW)] < 0$ for any $1 < i < K+1$ and $i \notin \mathcal{Z}$. Note that by using a similar argument, we have $\mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_1)(c_1 - b_1W)] > 0$ and $\mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_{K+1})(c_K - b_KW)] < 0$.

Next, we show that for any $i \in \mathcal{Z}$, we can construct a policy such that the IC constraints are strictly satisfied. Consider a given $i \in \mathcal{Z}$. Suppose $n_i = \lfloor c_{i-1}/b_{i-1}\rfloor = \lfloor c_i/b_i\rfloor$, $y_{i-1} = c_{i-1}/b_{i-1} - n_i$ and $1 > y_i = c_i/b_i - n_i$. Note that $y_i > y_{i-1} \geq 0$.

We construct a new policy $\tilde{\boldsymbol{\sigma}}$ as follows.

$$\tilde{\sigma}^l(S_j|n_i) = \begin{cases} \epsilon_0^l & \text{for } j = i \\ (1-\epsilon_0^l)\sigma^l(S_j|n) & \text{for } j \neq i \end{cases}, \qquad \tilde{\sigma}^l(S_j|n_i+1) = \begin{cases} \epsilon_1^l & \text{for } j = i \\ (1-\epsilon_1^l)\sigma^l(S_j|n) & \text{for } j \neq i \end{cases},$$

and for all other n, $\tilde{\sigma}^l(S_j|n) = \sigma^l(S_j|n)$ for all $j \in [K+1]$, where $\boldsymbol{\sigma}$ is the monotone policy corresponding to $\{\mathcal{T}_i\}_{i\in[K+1]}$ defined above. Then, using similar argument as in Eq. (27), we have that the IC constraints are satisfied for all $j \notin \mathcal{Z}$.

We argue that there exists $\{\epsilon_0^l\}_{l\in[L]}$ and $\{\epsilon_1^l\}_{l\in[L]}$ such that the IC constraints are strictly satisfied for the given i. Suppose $\tilde{\boldsymbol{\Pi}}$ is the corresponding steady-state distribution tuple. Then, for the IC constraint to be strictly satisfied, we need that

$$\begin{aligned} 0 &< \sum_{n\geq 0} \tilde{\pi}_n^l \tilde{\sigma}^l(S_i|n)(c_i - b_i n) \\ &= \tilde{\pi}_{n_i}^l \epsilon_0^l (c_i - b_i n_i) + \tilde{\pi}_{n_i+1}^l \epsilon_1^l \big(c_i - b_i(n_i+1)\big) \overset{(a)}{=} \tilde{\pi}_{n_i}^l \epsilon_0^l b_i y_i - \tilde{\pi}_{n_i+1}^l \epsilon_1^l b_i (1-y_i), \end{aligned}$$

where (a) follows using $c_i = b_i(n_i + y_i)$. Thus, we need that, $\tilde{\pi}_{n_i}^l \epsilon_0^l y_i > \tilde{\pi}_{n_i+1}^l \epsilon_1^l (1-y_i)$. Similarly, to satisfy $\sum_{n\geq 0} \tilde{\pi}_n^l \tilde{\sigma}^l(S_i|n)(c_{i-1} - b_{i-1} n) < 0$, we need that $\tilde{\pi}_{n_i}^l \epsilon_0^l y_{i-1} < \tilde{\pi}_{n_i+1}^l \epsilon_1^l (1-y_{i-1})$. Further,

$$\frac{1}{\tilde{\pi}_{n_i}^l} \tilde{\pi}_{n_i+1}^l = \lambda_l P_i \epsilon_0^l + (1-\epsilon_0^l) \sum_{j\neq i} \lambda_l P_j \sigma^l(S_j|n_i).$$

As such $\tilde{\pi}_{n_i+1}^l / \tilde{\pi}_{n_i}^l$ does not depend on ϵ_1^l. So, for all $l \in [L]$, we can pick ϵ_0^l small enough, and choose

$$\epsilon_1^l = \frac{\epsilon_0^l \tilde{\pi}_{n_i}^l}{2\tilde{\pi}_{n_i+1}^l} \Big(\frac{y_{i-1}}{1-y_{i-1}} + \frac{y_i}{1-y_i} \Big),$$

such that both the IC constraints corresponding to the signal S_i are strictly satisfied.

As the above argument holds for any $i \in \mathcal{Z}$, we have a sequence of policies $\{\boldsymbol{\sigma}\} \cup \{\boldsymbol{\sigma}_{(i)}\}_{i\in\mathcal{Z}}$, such that: (i) For $\boldsymbol{\sigma}$, all the IC constraints are strictly satisfied for any S_j such that $j \notin Z$, and (ii) For $\boldsymbol{\sigma}_{(i)}$ where $i \in \mathcal{Z}$, IC constraints corresponding to the signal S_i are strictly satisfied. Suppose $(\mathbf{p}, \mathbf{w}) = \chi(\boldsymbol{\sigma})$ and $(\mathbf{p}_{(i)}, \mathbf{w}_{(i)}) = \chi(\boldsymbol{\sigma}_{(i)})$ and let $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$ is the convex combination of $(\mathbf{p}, \mathbf{w})$ and $(\mathbf{p}_{(i)}, \mathbf{w}_{(i)})$'s with uniform weight $1/(1+|\mathcal{Z}|)$. Then, for the policy $\hat{\boldsymbol{\sigma}}$ corresponding to $(\hat{\mathbf{p}}, \hat{\mathbf{w}})$, all the IC constraints are strictly satisfied. This shows that the Slater's condition holds.

<u>Strong Duality:</u>

Using the variables $(\mathbf{p}, \mathbf{w})$, where $p^l(S_i) = \mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_i)|\Lambda = \lambda_l]$ and $w^l(S_i) = \mathbb{E}_{\boldsymbol{\sigma}}[W\xi(S_i)|\Lambda = \lambda_l]$, we have that $\mathbb{E}_{\boldsymbol{\sigma}}[\xi(S_i)] = \sum_{l\in[L]} q_l p^l(S_i)$ and $\mathbb{E}_{\boldsymbol{\sigma}}[W\xi(S_i)] = \sum_{l\in[L]} q_l w^l(S_i)$, and so,

$$\begin{aligned} H(\boldsymbol{\sigma}; \boldsymbol{\tau}, \tilde{\boldsymbol{\tau}}) &= \sum_{l\in[L]} \sum_{i=1}^{K+1} q_l \big[\lambda_l p^l(S_i) + \tau_i (c_i p^l(S_i) - b_i w^l(S_i)) - \tilde{\tau}_i (c_{i-1} p^l(S_i) - b_{i-1} w^l(S_i)) \big] \\ &= \sum_{l\in[L]} \sum_{i=1}^{K+1} q_l \big[\big(\lambda_l + \tau_i c_i - \tilde{\tau}_i c_{i-1} \big) p^l(S_i) - \big(\tau_i b_i - \tilde{\tau}_i b_{i-1} \big) w^l(S_i) \big], \end{aligned}$$

where $\tilde{\tau}_1 = c_0 = b_0 = 0$ and also, $\tau_{K+1} = c_{K+1} = b_{K+1} = 0$. From the above equation, we can define a function $\tilde{H}(\mathbf{p}, \mathbf{w}; \boldsymbol{\tau}, \tilde{\boldsymbol{\tau}})$, where $\tilde{H}(\mathbf{p}, \mathbf{w}; \boldsymbol{\tau}, \tilde{\boldsymbol{\tau}}) = H(\boldsymbol{\sigma}; \boldsymbol{\tau}, \tilde{\boldsymbol{\tau}})$ whenever $(\mathbf{p}, \mathbf{w}) = \xi(\boldsymbol{\sigma})$. Further, $\tilde{H}(\mathbf{p}, \mathbf{w}; \boldsymbol{\tau}, \tilde{\boldsymbol{\tau}})$ is affine in $(\mathbf{p}, \mathbf{w})$ and also in $(\boldsymbol{\tau}, \tilde{\boldsymbol{\tau}})$. Now, as the Slater's condition is satisfied and $\mathcal{X}$ is a compact and convex set, we use the Minimax theorem (see Willem 2012) to imply that

$$\max_{(\mathbf{p},\mathbf{w})\in\mathcal{X}} \min_{(\boldsymbol{\tau},\tilde{\boldsymbol{\tau}})\geq \mathbf{0}} \tilde{H}(\mathbf{p}, \mathbf{w}; \boldsymbol{\tau}, \tilde{\boldsymbol{\tau}}) = \min_{(\boldsymbol{\tau},\tilde{\boldsymbol{\tau}})\geq \mathbf{0}} \max_{(\mathbf{p},\mathbf{w})\in\mathcal{X}} \tilde{H}(\mathbf{p}, \mathbf{w}; \boldsymbol{\tau}, \tilde{\boldsymbol{\tau}}).$$

This shows that strong duality holds. Now, the KKT conditions in Eq. (8) follows simply as a consequence of the Minimax theorem.

$\underline{\tau_i^* \tilde{\tau}_i^* = 0}$:

In the final part, we show that $\tau_i^* \tilde{\tau}_i^* = 0$. Note that, we do not have the variables τ_1^* or $\tilde{\tau}_{K+1}^*$, and so, with slight abuse of notation, we choose $\tau_1^* = \tilde{\tau}_{K+1}^* = 0$. For any $i \in [K]$ and $i \neq 1$ such that $\mathbb{E}_{\boldsymbol{\sigma}^*}\left[\xi(S_i)\right] > 0$, the constraint $\mathbb{E}_{\boldsymbol{\sigma}^*}\left[\xi(S_i) h_{i-1}(W)\right] = 0$ implies that $\mathbb{E}_{\boldsymbol{\sigma}^*}\left[W \xi(S_i)\right] = \frac{c_{i-1}}{b_{i-1}} \mathbb{E}_{\boldsymbol{\sigma}^*}\left[\xi(S_i)\right]$, and similarly, the constraint $\mathbb{E}_{\boldsymbol{\sigma}^*}\left[\xi(S_i) h_i(W)\right] = 0$ implies that $\mathbb{E}_{\boldsymbol{\sigma}^*}\left[W\xi(S_i)\right] = \frac{c_i}{b_i} \mathbb{E}_{\boldsymbol{\sigma}^*}\left[\xi(S_i)\right]$. As such, we cannot have $\mathbb{E}_{\boldsymbol{\sigma}^*}\left[\xi(S_i) h_{i-1}(W)\right] = \mathbb{E}_{\boldsymbol{\sigma}^*}\left[\xi(S_i) h_i(W)\right] = 0$, and so, by complementary slackness, we get $\tau_i^* \tilde{\tau}_i^* = 0$.

E.2. Proof of Lemma 3

For ease of notation, consider a fixed arrival rate $\lambda = \lambda_l$, and drop the subscript (and superscript) corresponding to arrival rate index. Recall that $\beta(n-1) = \max_{i\in[K+1]}\{\lambda P_i \beta(n) + H_i(n)\} - b^*$, where

$$H_i(n) = \begin{cases} \lambda P_1 + \tau_1 h_1(n) & \text{for } i = 1 \\ \lambda P_i + \tau_i h_i(n) - \tilde{\tau}_i h_{i-1}(n) & \text{for } 1 < i < K+1 \\ -\tilde{\tau}_{K+1} h_K(n) & \text{for } i = K+1. \end{cases}$$

For any $n \geq 1$ and j such that $\sigma(S_j | n+1) > 0$, we have

$$\begin{aligned} \Delta^2 \beta(n-1) &= \beta(n+1) + \beta(n-1) - 2\beta(n) \\ &= \max_{i\in[K+1]} \Big[\lambda P_i \beta(n+2) + H_i(n+2)\Big] + \max_{i\in[K+1]} \Big[\lambda P_i \beta(n) + H_i(n)\Big] \\ &\quad - 2 \max_{i\in[K+1]} \Big[\lambda P_i \beta(n+1) + H_i(n+1)\Big] \\ &\overset{(a)}{\geq} \Big[\lambda P_j \beta(n+2) + H_j(n+2)\Big] + \Big[\lambda P_j \beta(n) + H_j(n)\Big] \\ &\quad - 2 \max_{i\in[K+1]} \Big[\lambda P_i \beta(n+1) + H_i(n+1)\Big] \\ &\overset{(b)}{=} \Big[\lambda P_j \Delta^2 \beta(n) + \Delta^2 H_j(n)\Big] \\ &\overset{(c)}{=} \lambda P_j \Delta^2 \beta(n), \end{aligned} \tag{28}$$

where (a) holds for any $j \in [K+1]$, and so it holds for j such that $\sigma(S_j|n+1) > 0$; (b) follows as $\sigma(S_j|n+1) > 0$ only if $j \in \arg\max_{i\in[K+1]} \Big[\lambda P_i \beta(n+1) + H_i(n+1)\Big]$; and (c) follows using $\Delta^2 H_j(n) = 0$, as $H_j(n)$ is linear for all $j \in [K+1]$. As the above holds for all j such that $\sigma(S_j|n+1) > 0$, we have

$$\Delta^2\beta(n-1) \geq \lambda\Delta^2\beta(n)\sum_{i=1}^{K} P_i\sigma(S_i|n+1). \tag{29}$$

From above argument, if $\Delta^2\beta(n') \geq 0$ for any n', then for all $n < n'$, we have $\Delta^2\beta(n) \geq 0$.

Consider the first case where $\sum_{i=1}^K P_i\sigma(S_i|n'+1) = 0$ (i.e., $\sigma(S_{K+1}|n'+1) = 1$) for some n'. Then we trivially have that $\Delta^2\beta(n-1) \geq 0$ for all $n < n'$. Also, without loss of generality, if $\sigma(S_{K+1}|n'+1) = 1$ for some n', then we can choose $\sigma(S_{K+1}|n+1) = 1$ for all $n \geq n'$. This implies that $\Delta^2\beta(n-1) \geq 0$ for all $n \geq n'$ also.

Now, in the second case, suppose $\sum_{i=1}^K P_i\sigma(S_i|n+1) > 0$ for all n. Let $\Delta^2\beta(n-1) = -\epsilon < 0$ for some n. As $\sum_{i=1}^K P_i\sigma(S_i|n) > 0$ for all n, by using the flow balance equation, i.e., $\pi_{n+1} = \pi_n \sum_{i=1}^K P_i\sigma(S_i|n)$, we also have $\pi_n > 0$ for all n. Now from Eq. (29), we get that $\pi_{n+1}\Delta^2\beta(n-1) \geq \pi_{n+2}\Delta^2\beta(n)$. By recursively using this inequality, for all $t \geq 0$, we have

$$\pi_{n+1}\Delta^2\beta(n-1) \geq \pi_{n+t+1}\Delta^2\beta(n+t) \implies \Delta^2\beta(n+t) \leq -\epsilon\frac{\pi_{n+1}}{\pi_{n+t+1}}$$

Then, as $\lim_{t\to\infty}\pi_{n+t+1} = 0$, we have $\lim_{t\to\infty}\Delta^2\beta(n+t) = -\infty$. This in turn implies that $\lim_{t\to\infty}\beta(n+t) = -\infty$, violating the fact that for all $n \geq 1$,

$$\beta(n-1) \geq \lambda P_{K+1}\beta(n) + H_{K+1}(n) - b^* = \tilde{\tau}^*_{K+1}(b_K n - c_K) - b^*.$$

This creates a contradiction and so we have that $\Delta^2\beta(n) \geq 0$ for all n. As the above argument holds for any $l \in [L]$, it completes the proof of Lemma 3.

E.3. Proof of Lemma 4

For ease of notation, we consider a given arrival rate is $\lambda = \lambda_l$, and drop the superscript and subscript corresponding to arrival rate index. Suppose $\sigma(S_i|n) > 0$ for $n = \underline{n}_i$ and $n = \overline{n}_i$, where $\overline{n}_i > \underline{n}_i$. We need to show that $\sigma(S_{i^-}|n) = 0$ for all $\underline{n}_i < n < \overline{n}_i$ and $i^- < i$.

The argument holds trivially if $\overline{n}_i = \underline{n}_i + 1$, simply because there does not exist an n such that $\underline{n}_i < n < \overline{n}_i$. As such, we consider that $\overline{n}_i \geq \underline{n}_i + 2$.

Now, consider a fixed i^- with $i^- < i$. We define the difference in the value function of two signal S_i and S_{i^-} as $Q(i,i^-;n) := Q(n,S_i) - Q(n,S_{i^-})$. Recall that

$$Q(i,i^-;n) = \lambda(P_i - P_{i^-})\beta(n) + H_i(n) - H_{i^-}(n), \tag{30}$$

where

$$H_i(n) = \begin{cases} \lambda P_1 + \tau_1 h_1(n) & \text{for } i = 1 \\ \lambda P_i + \tau_i h_i(n) - \tilde{\tau}_i h_{i-1}(n) & \text{for } 1 < i < K+1 \\ -\tilde{\tau}_{K+1} h_K(n) & \text{for } i = K+1. \end{cases}$$

We know that $\sigma(S_i|n) > 0$ only if $Q(n, S_i) = \max_j\{Q(n, S_j)\}$, and so, $Q(i, i^-; n) \geq 0$ for $n = \underline{n}_i$ and $n = \overline{n}_i$. In order to show that $\sigma(S_{i^-}|n) = 0$ for all $\underline{n}_i < n < \overline{n}_i$, it suffices to show that $Q(i, i^-; n) > 0$ for all $\underline{n}_i < n < \overline{n}_i$.

We recall some properties of $\beta(\cdot)$ function. From Lemma 3, we have that $\Delta^2\beta(n) \geq 0$. Suppose

$$N_0 := \min\{n : \Delta^2\beta(n) = 0\}.$$

From Lemma 3, we get that $\Delta^2\beta(n) = 0$ for all $n \geq N_0$ and $\Delta^2\beta(n) > 0$ for all $n < N_0$ by the definition of N_0. This implies that $\beta(n)$ is linear for $n > N_0$, and so there exists Δ_∞ such that $\beta(n) = \beta(N_0) + (n - N_0)\Delta_\infty$.

Now, we prove the result using a case-by-case analysis.

Case 1: [$\underline{n}_i < N_0$] In this case, for any $i^- < i$, using $H_i(n)$ and $H_{i^-}(n)$ are linear for all n, we have

$$\Delta^2 Q(i, i^-; \underline{n}_i) = \lambda(P_i - P_{i^-})\Delta^2\beta(\underline{n}_i) < 0,$$

where the last inequality follows as $P_{i^-} > P_i$ for any $i^- < i$, and $\Delta^2\beta(\underline{n}_i) > 0$ as $\underline{n}_i < N_0$. Further, we also get, for any n, $\Delta^2 Q(i, i^-; n) = \lambda(P_i - P_{i^-})\Delta^2\beta(n) \leq 0$ as $\Delta^2\beta(n) \geq 0$ for all n.

Now, we have three properties: (i) $Q(i, i^-; n) \geq 0$ for $n \in \{\underline{n}_i, \overline{n}_i\}$, (ii) $\Delta^2 Q(i, i^-; n) \leq 0$ for all $\underline{n}_i < n < \overline{n}_i$, and (iii) $\Delta^2 Q(i, i^-; n) < 0$ for $n = \underline{n}_i$. From these three properties, we get that $Q(i, i^-; n) > 0$ for all $\underline{n}_i < n < \overline{n}_i$ and $i^- < i$.

Case 2: [$\underline{n}_i \geq N_0$] In this case, we have that $\Delta^2 Q(i, i^-; n) = \lambda(P_i - P_{i^-})\Delta^2\beta(n) = 0$ for all $n \geq \underline{n}_i \geq N_0$. As such, $Q(i, i^-; n)$ is linear in n for $n \geq \underline{n}_i$.

- Case 2.1: [$Q(i, i^-; n) > 0$ at $n = \underline{n}_i$] In this case, as $Q(i, i^-; n)$ is linear in n for $n \geq \underline{n}_i$, and $Q(i, i^-; n) > 0$ for $n = \underline{n}_i$ and $Q(i, i^-; n) \geq 0$ for $n = \overline{n}_i$, we get $Q(i, i^-; n) > 0$ for all $\underline{n}_i \leq n < \overline{n}_i$.
- Case 2.2: [$Q(i, i^-; n) > 0$ at $n = \overline{n}_i$] This case is similar to Case 2.1.

Case 2.3: [$\underline{n}_i \geq N_0$ and $Q(i, i^-; n) = 0$ at $n = \underline{n}_i$ and $n = \overline{n}_i$] In this case, as $Q(i, i^-; n) = 0$ at two distinct point, and as already mentioned, $Q(i, i^-; n)$ is linear for $n \geq \underline{n}_i$. Thus, $Q(i, i^-; n) = 0$ for all $n \geq \underline{n}_i$. We argue that this leads to a contradiction, i.e., we cannot have $Q(i, i^-; n) = 0$ for all $n \geq \underline{n}_i$. Before providing the contradiction argument, we provide three intermediary steps:

- First, we argue that if $i \in \arg\max_j\{Q(S_j,n)\}$ for some $n \geq N_0+2$, then, $\beta(n-1) = \lambda P_i\beta(n) + H_i(n) - b^*$ for all $n \geq N_0+1$. This can be seen as follows. Recall that $\beta(\cdot)$ is linear for $n \geq N_0$, and so we denote $\Delta\beta(n) = \Delta_\infty$ for $n \geq N_0$. Suppose $i \in \arg\max_j\{Q(S_j,n)\}$ for some $n \geq N_0+2$, then $Q(S_i,n) = \max_j\{Q(S_j,n)\} = \beta(n-1)$.

  As $\beta(\cdot)$ is linear, we have that $Q(S_i,n)$ is also linear for $n \geq N_0$. If $\Delta Q(S_i,n) > \Delta_\infty$, then, $Q(S_i,n+1) > \beta(n)$, which violates the fact that $\beta(n) = \max_j\{Q(S_j,n+1)\}$. Similarly, if $\Delta Q(S_i,n) < \Delta_\infty$, then, $Q(S_i,n-1) > \beta(n-2)$ which violates the fact that $\beta(n-2) = \max_j\{Q(S_j,n-1)\}$. As such $\Delta Q(S_i,n) = \Delta_\infty$. Further, as $Q(S_i,n) = \beta(n-1)$ for $n \geq N_0 + 2$, where both $Q(S_i,n)$ and $\beta(n-1)$ are linear with same slope (Δ_∞), we have that $\beta(n-1) = Q(S_i,n) = \lambda P_i\beta(n) + H_i(n) - b^*$ for all $n \geq N_0+1$.
- Next, we argue that $\Delta_\infty \geq 0$. This holds simply because $\beta(n) = \max_j\{Q(n,S_j)\} \geq Q(n,S_{K+1}) = \tilde{\tau}_{K+1}(b_K n - c_K) - b^*$, where $\tilde{\tau}_{K+1} \geq 0$. If $\Delta_\infty < 0$, then there exists n for which $\beta(n) < \tilde{\tau}_{K+1}(b_K n - c_K) - b^*$, which contradicts the previous statement. Thus, $\Delta_\infty \geq 0$.
- Now, we argue that $\lambda P_i \leq 1$ and $\lambda P_{i^-} \leq 1$. Note that as $\beta(n-1) = \lambda P_i\beta(n) + H_i(n) - b^*$ for all $n \geq N_0+1$, $i \in \arg\max_j\{\lambda P_j\beta(n) + H_j(n)\}$ and so we can construct a different policy $\tilde{\boldsymbol{\sigma}}$ such that $\tilde{\sigma}(S_i|n) = 1$ for all $n \geq N_0+1$. As this new policy $\tilde{\boldsymbol{\sigma}}$ is also a solution of the Bellman equation, it should hold that the overall reward under $\tilde{\boldsymbol{\sigma}}$ matches the overall reward under the policy $\boldsymbol{\sigma}$ (where $\boldsymbol{\sigma}$ is optimal). However, if $\lambda P_i > 1$, we would have that the steady state distribution does not exist under the policy $\tilde{\boldsymbol{\sigma}}$, and so the overall reward under $\tilde{\boldsymbol{\sigma}}$ and $\boldsymbol{\sigma}$ cannot match. This creates a contradiction and so we have $\lambda P_i \leq 1$. Similar argument holds i^-, which gives us $\lambda P_{i^-} \leq 1$.

Now, we provide the arguments for this case. As $\sigma(S_i|n) > 0$ at $n = \overline{n}_i$ and $Q(i,i^-;n) = 0$ at $n = \overline{n}_i$, we have $i, i^- \in \arg\max_j\{Q(S_j,\overline{n}_i)\}$ where $\overline{n}_i \geq \underline{n}_i + 2 \geq N_0+2$. Now, by using the previous argument, $\beta(n) = \lambda P_i\beta(n+1) + H_i(n+1) - b^*$ for $n \geq N_0$. Note that this implies

$$\Delta_\infty = \lambda P_i\Delta_\infty + \Delta H_i \implies \Delta H_i = (1-\lambda P_i)\Delta_\infty, \tag{31}$$

where $\Delta H_i := \Delta H_i(n)$ for all n as $H_i(n)$ is linear. Similarly, we also have $\Delta H_{i^-} = (1-\lambda P_{i^-})\Delta_\infty$. By using $Q(i,i^-;n) = \lambda(P_i - P_{i^-})\beta(n) + H_i(n) - H_{i^-}(n)$, we have that

$$\begin{aligned}
&Q(i,i^-;n) - \lambda P_i Q(i,i^-;n+1)\\
&\quad= \lambda(P_i - P_{i^-})\big(\beta(n) - \lambda P_i\beta(n+1)\big) + H_i(n) - H_{i^-}(n) - \lambda P_i\big(H_i(n+1) - H_{i^-}(n+1)\big)\\
&\quad\overset{(a)}{=} \lambda(P_i - P_{i^-})\big(H_i(n+1) - b^*\big) + H_i(n) - \lambda P_i H_i(n+1) - H_{i^-}(n) + \lambda P_i H_{i^-}(n+1)\\
&\quad= -\lambda(P_i - P_{i^-})b^* + H_i(n) - \lambda P_{i^-}H_i(n+1) - H_{i^-}(n) + \lambda P_i H_{i^-}(n+1),
\end{aligned} \tag{32}$$

where (a) follows using $\beta(n) = \lambda P_i\beta(n+1) + H_i(n+1) - b^*$ for $n \geq N_0$. Also,

$$\begin{aligned}
H_i(n) &= \lambda P_i + \tau_i(c_i - b_i n) + \tilde{\tau}_i(b_{i-1}n - c_{i-1})\\
&\quad= \lambda P_i + (n-\chi_i)\Delta H_i = \lambda P_i + (1-\lambda P_i)(n-\chi_i)\Delta_\infty,
\end{aligned}$$

where $\chi_i := \frac{c_i}{b_i}\mathbb{1}_{\{\tau_i\neq 0\}} + \frac{c_{i-1}}{b_{i-1}}\mathbb{1}_{\{\tilde{\tau}_i\neq 0\}}$, where we use that fact that $\tau_i\tilde{\tau}_i = 0$ as mentioned in Lemma 2. Thus,

$$H_i(n) - \lambda P_{i^-} H_i(n+1) = \big(1-\lambda P_{i^-}\big)\big(\lambda P_i + (1-\lambda P_i)(n-\chi_i)\Delta_\infty\big) - \lambda P_{i^-}(1-\lambda P_i)\Delta_\infty$$

The above equality also holds for i^-. Thus, by substituting this in Eq. (32), we have

$$\begin{aligned} &Q(i,i^-;n) - \lambda P_i Q(i,i^-;n+1) \\ &\quad = \lambda(P_i - P_{i^-})(1+\Delta_\infty - b^*) + (1-\lambda P_i)(1-\lambda P_{i^-})\Delta_\infty(\chi_{i^-} - \chi_i) \\ &\quad \le \lambda(P_i - P_{i^-})(1-b^*) \end{aligned}$$

where the last inequality holds by using $\lambda P_i \le 1$, $\lambda P_{i^-} \le 1$ and $\chi_{i^-} \le \chi_i$, as c_i/b_i follow an ascending order. Finally, as b^* denotes the optimal throughput for the throughput optimization problem in (4), as such, $b^* < 1$. This implies that $\lambda(P_i - P_{i^-})(1-b^*) < 0$ and so $Q(i,i^-;n) - \lambda P_i Q(i,i^-;n+1) < 0$. Thus, we cannot have $Q(i,i^-;n) = Q(i,i^-;n+1) = 0$. This creates a contradiction that $Q(i,i^-;n) = 0$ for all $n \ge \underline{n}_i$. This completes the proof.

As the above argument holds for any $l \in [L]$, it completes the proof of

Appendix F: Proof of Proposition 6

In order to prove Proposition 6, we use a similar methodology as in the proof of Theorem 1, as provided in Section 6. We consider $\boldsymbol{\sigma}$ to be an optimal policy, and suppose b^* is the corresponding optimal value, i.e., $b^* = \max_{\boldsymbol{\sigma}}\ H(\boldsymbol{\sigma})$, where

$$H(\boldsymbol{\sigma}) = \mathbb{E}_{\boldsymbol{\sigma}}\Big[\sum_{i=1}^{K}\xi(S_i)H_i(W)\Big],$$

where $H_i(n) = \alpha_i - \gamma_i n$ Without loss of generality, we assume that if $\sigma(S_{K+1}|n) = 1$ for some n, then $\sigma(S_{K+1}|n') = 1$ for all $n' \ge n$.

First, we assume that, for the optimal policy $\boldsymbol{\sigma}$, there exists n_0 and ϵ such that $\sum_{i=1}^{K}\lambda P_i\sigma(S_i|n) < 1-\epsilon$ for all $n \ge n_0$ and for all $l \in [L]$. Under this assumption, we can write the Bellman optimality equation as follows.

By using similar argument as for Eq. (11), for any randomized optimal policy, $\sigma(S_i|n) > 0$ only if $i \in \arg\max_{i\in[K+1]}\Big\{H_i(n) + \lambda P_i \Delta V(n)\Big\}$, where $V(n)$ denotes the differential cost function for queue length n. And by denoting $\beta(n) = \Delta V(n)$, we have that for all $n \ge 1$,

$$\beta(n-1) = \max_{i\in[K+1]}\big\{\lambda P_i\beta(n) + H_i(n)\big\} - b^*, \tag{33}$$

and at $n = 0$, we have $b^* = \max_{i\in[K+1]}\big\{\lambda P_i\beta(0) + H_i(0)\big\}$. Next we present some properties of $\beta(\cdot)$ function.

By using the same argument as in the proof of Lemma 3, for any $n \geq 1$ and any $i \in [K+1]$ such that $\sigma(S_i|n+1) > 0$,

$$\Delta^2\beta(n-1) \geq \lambda P_i \Delta^2\beta(n), \tag{34}$$

and $\beta(n)$ is weakly convex, i.e., $\Delta^2\beta(n) \geq 0$ for all $n \geq 0$. The above equation implies that if $\beta(\cdot)$ is strictly convex for any n, i.e., $\Delta^2\beta(n) > 0$, then for all $n' < n$, we have $\Delta^2\beta(n') > 0$. Suppose

$$N_0 := \min\{n : \Delta^2\beta(n) = 0\}.$$

From Eq. (34), we get that $\Delta^2\beta(n) = 0$ for all $n \geq N_0$. This implies that $\beta(n)$ is linear for $n > N_0$, and so there exists Δ_∞ such that $\beta(n) = \beta(N_0) + (n - N_0)\Delta_\infty$.

<u>If $N_0 < \infty$, then $\Delta\beta(n) = 0$ for all $n \geq N_0$.</u>

Note that if $\Delta_\infty < 0$, then there exists n large enough such that $\beta(n) < -b^*$. However, $\beta(n) = \max_{j\in[K+1]}\Big[\lambda P_j\beta(n+1) + H_j(n+1)\Big] - b^* \geq \lambda P_{K+1}\beta(n+1) + H_{K+1}(n+1) - b^*$. As such we have $\Delta_\infty \geq 0$.

First, we consider the case where there exists $n \geq 1$ such that $\sigma(S_{K+1}|n) = 1$. Then without loss of generality, we have $\sigma(S_{K+1}|n') = 1$ for all $n' > n$. Also, as $\beta(n') = \lambda P_j\beta(n'+1) + H_j(n'+1) - b^*$ for any j such that $\sigma(S_j|n') > 0$, we have that $\beta(n') = -b^*$ for all $n' \geq n$, using the fact that $\sigma(S_{K+1}|n') = 1$ for all $n' > n$, $P_{K+1} = 0$ and $H_{K+1}(\cdot) = 0$. This implies that $\Delta\beta(n') = 0$ for all $n' > n$ and so $\Delta_\infty = 0$ in this case.

Next, we consider the case where $\sigma(S_{K+1}|n) < 1$ for all n. This implies that the steady state distribution satisfies $\pi_n > 0$ for all n. For any $n \geq 1$, we have that

$$\begin{aligned}
\Delta\beta(n-1) &= \beta(n) - \beta(n-1) \\
&= \max_{j\in[K+1]}\Big[\lambda P_j\beta(n+1) + H_j(n+1)\Big] - \max_{j\in[K+1]}\Big[\lambda P_j\beta(n) + H_j(n)\Big] \\
&\overset{(a)}{=} \Big[\lambda P_i\beta(n+1) + H_i(n+1)\Big] - \max_{j\in[K+1]}\Big[\lambda P_j\beta(n) + H_j(n)\Big] \\
&\overset{(b)}{\leq} \Big[\lambda P_i\beta(n+1) + H_i(n+1)\Big] - \Big[\lambda P_i\beta(n) + H_i(n)\Big] \\
&= \lambda P_i\Delta\beta(n) + \Delta H_i(n) \\
&= \lambda P_i\Delta\beta(n) - \gamma_i
\end{aligned} \tag{35}$$

where (a) holds for any i such that $\sigma(S_i|n+1) > 0$; and (b) holds for any $i \in [K+1]$ and so it also holds for i such that $\sigma(S_i|n+1) > 0$. Now, as $\gamma_i \geq 0$ for all i, $\Delta\beta(n-1) \leq \lambda P_i\Delta\beta(n)$.

As the above equation holds for any i such that $\sigma(S_i|n+1) > 0$, we also get that

$$\Delta\beta(n-1) \leq \Delta\beta(n)\sum_{i=1}^{K+1}\lambda P_i\sigma(S_i|n+1).$$

By using the balance equation, i.e., $\pi_{n+2} = \pi_{n+1} \sum_{i=1}^{K+1} \lambda P_i \sigma(S_i|n+1)$, we have that for all $n \geq 1$,

$$\pi_{n+1}\Delta\beta(n-1) \leq \pi_{n+2}\Delta\beta(n) \leq \pi_{n+m+2}\Delta\beta(n+m).$$

Thus, for $n \geq N_0 + 1$

$$\Delta\beta(n+m) \geq \frac{\pi_{n+1}}{\pi_{n+m+2}}\Delta\beta(n-1) = \frac{\pi_{n+1}}{\pi_{n+m+2}}\Delta_\infty,$$

where $\lim_{m\to\infty} \frac{\pi_{n+1}}{\pi_{n+m+2}} = \infty$. Now, if $\Delta_\infty > 0$, it implies that there exists m such that $\Delta\beta(n+m) > \Delta_\infty$, which contradicts the fact that $\Delta\beta(n+m) = \Delta_\infty$ for $n \geq N_0 + 1$. As such $\Delta_\infty = 0$. □

Claim 2. *Suppose N_1 and N_2 are queue length values such that $\sigma(S_i|n) > 0$ for $n \in \{N_1, N_2\}$, then, for all $N_1 < n < N_2$, we have that $\sigma(S_j|n) = 0$ for all $j < i$.*

Proof of the Claim: We define the value function for S_i at queue length n to be $Q(n, S_i) := \lambda P_i \beta(n) + H_i(n) - b^*$. As such, from Eq. (33), $\beta(n-1) = \max_{i\in[K+1]} Q(n, S_i)$. And for any pair of signals (S_i, S_j), we have

$$Q(n, S_i) - Q(n, S_j) = \lambda(P_i - P_j)\beta(n) + H_i(n) - H_j(n). \tag{36}$$

Note that, from Eq. (33), at any $n \geq 0$, $\sigma(S_i|n) > 0$ only if $Q(n, S_i) = \max_j\{Q(n, S_j)\}$.

To prove the above claim, we use similar argument as in the proof of Lemma 4. Suppose $i, i^- \in [K+1]$ and $i^- < i$ and $\sigma(S_i|n) > 0$ for $n = \underline{n}_i$ and $n = \overline{n}_i$, where $\overline{n}_i > \underline{n}_i$. We need to show that $\sigma(S_{i^-}|n) = 0$ for all $\underline{n}_i < n < \overline{n}_i$ and $i^- < i$. The argument holds trivially if $\overline{n}_i = \underline{n}_i + 1$. As such, we consider that $\overline{n}_i \geq \underline{n}_i + 2$.

Suppose $Q(i, i^-; n) := Q(n, S_i) - Q(n, S_{i^-})$. We know that $\sigma(S_i|n) > 0$ only if $Q(n, S_i) = \max_j\{Q(n, S_j)\}$. It suffices to show that $Q(i, i^-; n) > 0$ for all $\underline{n}_i < n < \overline{n}_i$ and $i^- < i$.

Note that $Q(i, i^-; n) \geq 0$ for $n = \underline{n}_i$ and $n = \overline{n}_i$. Suppose $P_i = P_{i^-}$, then we have that

$$Q(i, i^-; n) = H_i(n) - H_{i^-}(n),$$

which is linear in n. Then, if $Q(i, i^-; n) = 0$ at both $n = \underline{n}_i$ and $n = \overline{n}_i$, then we have that $H_i(n) = H_{i^-}(n)$, for all n. This means that means that $\alpha_i = \alpha_{i^-}$ and $\gamma_i = \gamma_{i^-}$ which creates redundancy and violates the assumption of the model. As such, either at $n = \underline{n}_i$ or at $n = \overline{n}_i$, we must have $Q(i, i^-; n) > 0$, and in both the case $Q(i, i^-; n) > 0$ for $\underline{n}_i < n < \overline{n}_i$ (using $Q(i, i^-; n) \geq 0$ for $n = \underline{n}_i$ and $n = \overline{n}_i$)

So, now we consider the case that $P_i < P_{i^-}$.

- Case 1: [$\underline{n}_i < N_0$] This case is same as the Case 1 in the proof of Lemma 4.
- Case 2: [$\underline{n}_i \geq N_0$] By using Eq. (35), for any i such that $\sigma(S_i|n+2) > 0$, $\Delta\beta(n) \leq \lambda P_i \Delta\beta(n+1) - \gamma_i$, and so we get

$$\Delta\beta(n) \leq \Delta\beta(n+1) \sum_{i=1}^{K+1} \lambda P_i \sigma(S_i|n+2) - \sum_{i=1}^{K+1} \gamma_i \sigma(S_i|n+2).$$

For any $n \geq N_0 + 2$, we have that $\Delta\beta(n) = \Delta\beta(n) = \Delta_\infty = 0$. As such, for any $n \geq N_0 + 2$, $-\sum_{i=1}^{K+1} \gamma_i \sigma(S_i|n+2) \geq 0$. As $\gamma_i > 0$ for any $i \in [K]$, we get that $\sigma(S_i|n+2) = 0$ for all $i \in [K]$ and $n \geq N_0 + 2$. This implies that $\sigma(S_{K+1}|n) = 1$ for all $n \geq N_0 + 2$.

Recall our earlier consideration that $\overline{n}_i \geq \underline{n}_i + 2$ as the argument holds trivially if $\overline{n}_i = \underline{n}_i + 1$. Now, if $\underline{n}_i \geq N_0$ (by the assumption of this case), then $\sigma(S_i|n) > 0$ for $n = \overline{n}_i$ (where $\overline{n}_i \geq \underline{n}_i + 2 \geq N_0 + 2$) implies that $i = K+1$. This implies that $\beta(n) = -b^*$ for all $n \geq N_0$, as $\Delta\beta(n) = 0$ is constant for all $n \geq N_0$ and $\beta(n) = -b^*$ for $n = \overline{n}_i$. As $\beta(n) = -b^*$ for all $n \geq N_0$ we can simply choose $\sigma(S_{K+1}|n) = 1$ for all $n \geq N_0$ as it satisfies the Bellman equation. This directly implies that the claim is satisfied in this case ($\underline{n}_i \geq N_0$).

This completes the proof of the Claim. □

In the last step, we verify that our initial assumption, that there exists n_0 and ϵ such that $\sum_{i=1}^{K} \lambda P_i \sigma(S_i|n) < 1 - \epsilon$ for all $n \geq n_0$ holds true. As $\boldsymbol{\sigma}$ is laminar, from the statement in Lemma 1, we have that for every signal S_i, there exists a queue length value N_i (possibly infinite), such that $\sigma(S_i|n) = 0$ for all $n > N_i$. Note that we also conclude that, there is a unique signal S_i, such that $N_i = \infty$ and for any $j \neq i$, $N_j < \infty$, as otherwise, the statement in Lemma 1 will be violated. Let $n_0 = \max_{j \neq i} N_j$. Then, for every $n > n_0$, the arrival rate is λP_i, as $\sigma(S_i|n) = 1$ for all $n > n_0$. If $\lambda^l P_i \geq 1$, then the underlying Markov chain would be unstable. Consequently, $\lambda P_i < 1$. Now, we can simply pick $\epsilon = (1 - \lambda P_i)$. This shows that our initial assumption indeed holds. This finally proves the result in Proposition 6.