Structural Estimation of Intertemporal Externalities with Application in ICU Admissions

Yiwen Shen
Department of ISOM, HKUST, yiwenshen@ust.hk

Carri W. Chan
Decision, Risk, and Operations, Columbia Business School, cwchan@columbia.edu

Fanyin Zheng
Decision, Risk, and Operations, Columbia Business School, fanyin.zheng@columbia.edu

Gabriel Escobar
Division of Research, Kaiser Permanente, gabriel.escobar@kp.org

Problem definition: In many service systems, the system manager needs to balance between addressing the needs of current customers and ensuring the system's ability to serve future customers. Such balancing behavior is particularly important in capacity-constrained systems with heterogeneous service levels, in which the manager needs to decide which level of service to provide to the current customer, taking into account the intertemporal externalities of their decisions.

Methodology/results: We develop a dynamic discrete choice model to describe the decision-making process in a gatekeeper system with multiple classes of servers and customers. The discount factor in the model captures how much the decision-maker internalizes the intertemporal externalities of their customer routing decisions. In contrast to most empirical studies in the literature which use a pre-specified discount factor, we establish joint identification of the discount factor and the utility parameters from data. We then apply the model to empirically study the Intensive Care Unit (ICU) admission decisions for Emergency Department (ED) patients. Using a large hospitalization data set, we find that there is large heterogeneity in the estimated discount factors across hospitals. Via counterfactual simulations, we show that understanding the balancing behavior from data is important for hospitals to manage their ICU congestion.

Managerial implications: Our results suggest that it is important to understand how the decision-maker internalizes the intertemporal externalities from data. In addition, the balancing behavior regarding current customers and future available capacity provides a potential channel for improving system performance.

Key words: structural estimation, dynamic discrete choice model, empirical operations management, healthcare, intensive care unit

---

1. Introduction

In managing service systems, the decision-maker often has to carefully balance between providing immediate service to current customers versus the impact of such actions on the system's ability to serve future customers. Such trade-off is particularly relevant when the system is operating in a resource constrained environment. While providing immediate service may benefit current customers, it will utilize the limited resources in the service system; thus, such an action may hinder access to service for future, perhaps more valuable, customers. This introduces *intertemporal externalities* in the decision-making, i.e., the impact of the current decision on the future system state matters.

These intertemporal externalities often arise in gatekeeper systems with heterogeneous service levels. In these systems, the decision-maker acts as a gatekeeper who decides which level of service to provide to the arriving customer. The higher-level service generally provides more benefits to the customers, but is usually more capacity-constrained than the lower-level service. When making decisions, the gatekeeper needs to consider not only the current customer's type, but also the system's state after the decision which determines its ability to serve future customers. Such gatekeeper systems can be found in many real-world settings and application areas, such as call centers (Hathaway et al. 2023) and border security check systems (Zhang et al. 2011). In healthcare settings, the 'customers' can be interpreted as patients, who usually receive different levels of care. In maternity departments, midwives act as gatekeepers to specialists, but may also try to serve patients themselves (Freeman et al. 2016). In emergency departments, physicians may admit waiting patients to intensive care units or ordinary floor units (Kim et al. 2015).

We take a structural estimation approach to understand how the decision-maker balances the needs of current customers versus saving system capacity for future customers in such gatekeeper systems. We first develop a dynamic discrete choice model that describes the decision-making process. The model includes two service units: a first-tier service unit (FSU) which provides higher level of service but with limited capacity, and a second-tier service unit (SSU) that provides lower level of service but with ample capacity. Customers are categorized as the low-priority class and high-priority class, with the former more likely to benefit from the first-tier service. In each period, the system considers three actions for each customer arriving or waiting at the gatekeeper (GK): 1) send the customer to the FSU for service, 2) send to the SSU, or 3) keep the customer waiting at the GK. We assume the system chooses the action that maximizes its discounted utility over an infinite horizon. Our model incorporates stochastic arrivals and departures, external arrivals to the FSU, and idiosyncratic shocks to the utility of each action.

While the FSU can provide the highest level of service, its limited capacity means that sending a customer to it may hinder the service for future high priority customers. Thus, the system needs to balance the utility of current customer and its capacity to serve future customers. We quantify how the system balances such trade-off in practice by estimating the *intertemporal discount factor* in the model from observed data. The discount factor represents the *relative* weight of the future periods' utility in the system's objective function. Thus, a larger discount factor implies that the system accounts more for the impact of its decisions on future system state; while a smaller discount factor means the system is more focused on current customers.

In this work, we estimate the discount factor and the utility parameters *jointly* from data.This is in stark contrast to the most of empirical studies with dynamic discrete choice models, which assume the discount factor to be known and only estimate the utility parameters (e.g., Rust 1987, Ryan 2012, Mehta et al. 2017, and Emadi and Staats 2020). Moreover, the discount factor is generally set at a high level close to one (e.g., 0.975) without empirical support and formal justifications. However, the discount factors in dynamic discrete choice models can vary dramatically depending on the context of the problem, or as a consequence of the behavioral bias of decision makers (Frederick et al. 2002). We estimate the discount factor from data to understand how the decision-maker responds to intertemporal externalities in practice, which also facilitates accurate evaluations of how different shocks impact the system performance.

It is well known that the dynamic discrete choice model can not be identified from choice data without further conditions (see e.g., Rust 1994, Magnac and Thesmar 2002). The non-identification stems from the existence of observationally equivalent structures: multiple combinations of discount factors and utility parameters can lead to the same choice probabilities for all states (Rust 1994). We circumvent this difficulty by leveraging the approach in Komarova et al. (2018) and extending it to our dynamic discrete choice model with state-dependent action set. Much like other general econometric methodologies, applying the theoretical results in Komarova et al. (2018) to our specific model is a challenging task that must be executed carefully for it to be valid. For this purpose, we develop and justify the necessary components of the identification method tailored to our setting, including the modeling assumptions, estimation procedure, and discussion on action set and normalization choice in the estimation.

We apply our discrete choice model and identification results of a gatekeeper system to a suitable and important empirical setting: hospital's Intensive Care Unit (ICU) admission decisions for patients admitted via the Emergency Department (ED). ICUs are specialized inpatient units that provide the highest level of care for the most critically ill patients. They are expensive medical resources and often operate at high occupancy levels (Halpern and Pastores 2015). As the ICU provides the highest level of care, swift admission generally benefits the patients who need ICU care. However, given the constrained capacity and high occupancy of the ICU, this may restrict access to ICU care for future, perhaps more severe, patients. Such a trade-off introduces intertemporal externalities in the system gatekeeper's ICU admission decisions.

We use the dynamic discrete choice model to estimate the ICU admission decisions, with the discount factor reflecting how hospitals balance the tradeoff between current versus future patients. We divide ED patients into low- and high-severity classes based on their risk scores, capturing the heterogeneity in their needs for ICU care. For each ED patient, the hospitals can make three decisions: admit to the ICU, admit to non-ICU units (e.g., wards), or keep the patient waiting in the ED. In this setting, the ED plays the role of gatekeeper for the ICU, which represents the first-tier service unit (FSU) with the highest level of service with limited capacity. The non-ICU units provide lower-level of care with ample capacity, and are viewed as the second-tier service unit (SSU). Thus, the decision-making process for the ED patient fits into our gatekeeper's decision-making described above. In particular, under certain model assumptions, a larger discount factor means that the hospital tends to be more responsive to ICU congestion by adjusting the admission probabilities proactively.

We implement our structural model empirically on a large data set from a major US hospital network, including 21 hospitals with more than 300,000 hospitalizations. The discount factor and action utilities in the model are estimated jointly from the observed admission decisions. Given the ICU admission is a complex process involving multiple stakeholders and units, the decision maker (hospital) may not be fully aware of its balancing behavior in the actual admission process. Moreover, the current discounting level is not necessarily optimal for hospitals' long-term goals. We find that the identified discount factor on average is much lower than the levels usually assumed in the literature. As shown in the behavioral economics literature, such present-focused behaviors may be explained by the challenge in forecasting future payoff (Gabaix and Laibson 2017), complexity in evaluating the decision (Imas et al. 2022), and framing effects related to gain and loss (Faralla et al. 2017). These factors can be relevant in the ICU admission process due to its complex and composite nature.

Moreover, we reveal large heterogeneity in the estimated discount factors across individual hospitals, i.e., hospitals balance the trade-off between current patient versus future system capacity very differently. A comprehensive correlational analysis shows that the discount factor tends to be larger for hospitals with more congested ICUs and longer ED waiting time. The heterogeneity in discount factors reveals a novel aspect of the practice variation in healthcare (Corallo et al. 2014). Our estimation results suggest that it is important to identify discount factor from data in different empirical settings, instead of using a pre-assumed level. We show that our estimated model fits the data well in multiple aspects and support our main findings by comprehensive robustness checks. In addition, the heterogeneity in hospitals' balancing behaviors is consistent with the evidence from reduced-form analysis using a multinomial logistic model.

We note that our interpretation of the estimated discount factors relies on the assumption that the per-period utility of each action is unaffected by ICU occupancy. An alternative explanation—that the clinical benefit of ICU care declines as the unit becomes congested—could, in principle, generate similar observed behaviors. While we cannot entirely rule out this mechanism, we provide empirical evidence that suggests, in our data setting, ICU congestion is not a primary driver of patient outcomes. Additionally, our identification and interpretation of the discount factor hinge on the pre-specified model assumptions, a practice common in the dynamic discrete choice model literature. That said, some of these assumptions cannot be verified empirically and if they are violated, it would invalidate our results.

The main benefit of structural estimation is its ability to evaluate counterfactual policy changes, as it identifies the underlying primitives that are invariant to such changes (e.g., Ryan 2012). We use counterfactual analyses to quantify how the hospital's balancing behaviors impact its high ICU congestion state, which can have serious repercussions on patient flow and outcomes (see, e.g., Kc and Terwiesch 2012, Kim et al. 2015). We first show that the traditional capacity intervention of adding one ICU bed can significantly decrease the ICU congestion. In addition, such benefit is more significant for hospitals with larger discount factors and/or more congested ICUs. Next, we consider a behavioral intervention of increasing the discount factor in hospitals' decision making. We find that for some hospitals, such intervention can reduce ICU congestion with magnitudes comparable to the costly act of adding one ICU bed. In reality, hospitals' present-focused behavior may arise due to multiple behavioral biases or system constraints. Thus, hospitals may be able to achieve more system-focused behavior by addressing these factors (e.g., more education or timely update of system state). We further develop a data-driven, easy-to-implement heuristic policy. This provides a practical nudge that can achieve most of the benefits in reducing ICU congestion generated by increasing discount factor. The heuristic policy suggests that hospitals should be more responsive to their ICU state and adjust the admission decisions when ICU occupancy is very high.

The rest of the paper is organized as follows. We conclude this section with a brief literature review. Section 2 develops the main dynamic discrete choice model and establishes the identification of discount factor. Section 3 describes how we apply the model to the ICU admission decisions. Section 4 presents the main empirical results. Section 5 conducts counterfactual studies and reveals the practical relevance of hospitals' discounting behaviors. We conclude the paper and discuss future research directions in Section 6. The auxiliary results are provided in the electronic companion. The proofs and other numerical results are collected in the online supplement for interested readers.

1.1. Literature Review

Our work is related to the following streams of literature: structural estimation in operations management; identification of dynamic discrete choice models; and empirical healthcare operations management.

The structural estimation approach has been increasingly used in operations management to understand the behaviors of customers and servers in different empirical settings. Bray et al. (2019) develop a dynamic discrete choice model and show that ration gaming exists in a multiechelon supply chain. Using structural estimation from air-travel industry, Li et al. (2014) find that customers sometimes delay purchases strategically when they expect the price to fall. Akşin et al. (2013) and Yu et al. (2016) take the structural estimation approach to study caller's abandonment behaviors in call centers. Emadi and Staats (2020) find that attrition of agents at a management firm appears to be insensitive to salary. Structural models have also been used in healthcare operations management. Olivares et al. (2008) use a newsvendor model to study how a hospital balances the costs of reserving too much versus too little operating room (OR) capacity. Rath and Rajaram (2022) use a choice model to estimate costs associated with OR scheduling of anesthesiologists. In this study, we develop a dynamic discrete choice model to understand the discounting behavior in a gatekeeper system with heterogeneous service levels. The dynamic choice model we use is pioneered by Rust (1987). We estimate the discount factor and utility parameters jointly from data, and apply our model to empirically study the ICU admission decisions from a large hospital network.

In most empirical studies with a dynamic choice model, the discount factor is pre-assumed at a very high level, such as 0.975 (e.g., Bray et al. 2019, Emadi and Staats 2020). This is because the joint identification of the discount factor and utility parameters is technically very challenging (e.g., Rust 1994, Magnac and Thesmar 2002). It has been a fundamental topic in economics and has attracted increasing attention in recent years (see, e.g., Komarova et al. 2018, Abbring and Daljord 2020, Abbring et al. 2025, Hao et al. 2025).

In the literature, the identification of discount factor has to rely on certain forms of the payoff functions (e.g., Komarova et al. 2018, Hao et al. 2025) and/or specific exclusion restriction conditions (e.g., Magnac and Thesmar 2002, Abbring and Daljord 2020. It has been done in only several recent studies in very different settings, e.g., Dubé et al. (2014), De Groote and Verboven (2019), Ching and Osborne (2020), and Dubé et al. (2024). They show that correctly identifying the discount factor is important for understanding agents' behaviors and evaluating policy impacts. In these cases, identification is achieved by leveraging special features in the data based on their respective empirical settings. To our best knowledge, we are the first to estimate a dynamic structural model with unknown discount factor in the empirical operations and healthcare management literature. Our identification and interpretation rely on pre-specified model assumptions (e.g., a linear utility function and constant utility parameters), which are explicitly stated and discussed in the paper.

Our empirical study of the ICU admission decisions contributes to the literature on empirical healthcare operations management. Various works have examined the behaviors and implications of gatekeepers in healthcare settings. Batt and Terwiesch (2016) find that, when the ED is congested, nurses will initiate diagnostic tests at the triage stage to reduce waiting times. Freeman et al. (2016) show that midwives in delivery units ration resource-intensive service and increase the rate of specialist referrals when workload increases. Freeman et al. (2021) find that adding a second gatekeeping stage in the ED reduces the rates of unnecessary hospitalization and wrongful patient discharge. Our work is also related to the studies on patient admission decisions and their implications. For instance, Edbrooke et al. (2011) and Kim et al. (2015) investigate the impact of ICU admissions on multiple patient outcomes. Chan et al. (2016) find that delays in ICU admission can increase ICU LOS, which can create more congestion in an already busy unit. When bed availability in the primary unit is limited, patients are typically rerouted to alternative units designated for a different service. Song et al. (2020) and Dong et al. (2018) show that such off-placement has important clinical and operational implications. There is substantial evidence that system state can influence physician admission decisions and patient flow dynamics. For instance, ICU congestion can impact who is admitted to the ICU (Kim et al. 2015) and when patients are discharged (Kc and Terwiesch 2012). In a setting similar to ours, Kim et al. (2020) propose a behavioral model and use controlled experiments to identify a number of factors which can bias physician's ICU admission decisions. In our work, we take a structural estimation approach to understand the balancing behaviors in ICU admissions decisions from data.

2. Model Set-Up and Identification of Discount Factor

In this section, we introduce the structural model of the gatekeeper system, explain the identification of the discount factor, and describe the estimation procedure.

2.1. Dynamic Discrete Choice Model for Gatekeeper Systems

In our structural model, we consider the routing decision of a gatekeeper (GK) for each customer about the downstream unit which serves them. We assume that there are two units, i.e., groups of servers, that provide different levels of service: the first-tier service unit (FSU) and the second-tier service unit (SSU). The FSU provides a higher level of service that generally benefits the customers. However, its capacity is usually more constrained than the SSU. There are three key features of the model. First, in each period, the GK considers which service unit, if any, to allocate to each customer waiting in the system. Second, such routing decision depends on both the customer type and the current system status, i.e., the remaining capacity of the service units. Finally, the model allows for the system to potentially consider the future state and to make dynamic decisions. We provide detailed descriptions of the model below.

We divide the arriving customers into two classes: the low-priority class and the high-priority class, represented by subscripts l and h, respectively. The model set-up can be extended to include multiple service units and customer classes as well. The customers in the high-priority class usually requires a higher level of service than those in the low-priority class. We assume that the GK has capacities Q_l and Q_h for the low-priority and high-priority classes, respectively. This can capture the limited waiting area at the GK. For $i \in \{l, h\}$, let $A_{i,t}$ be the number of class i customers arriving to the GK in time period t. $A_{i,t}$ follows a truncated Poisson distribution with rate $\lambda_{Q,i}$. While in theory $A_{i,t}$ can be unbounded, we truncate it at a sufficiently high upper bound $M_{A,i}$. This reduces the state space of the model to keep the estimation computationally feasible without imposing additional assumptions.

In addition to customers admitted via the GK, there are other arriving customers who may also require the first-tier service. Let E_t denote the number of customers arriving to the FSU via other channels in period t, referred to as the external arrivals. E_t is distributed according to a Poisson distribution with arrival rate λ_E. Every period, each current customer in the FSU completes their first-tier service and departs from the unit with probability μ_I. Thus, D_t, the number of customers departing from the FSU in each period, follows a binomial distribution. The total number of servers in the FSU is B, meaning that the first-tier service can be provided to at most B customers simultaneously. On the other hand, we assume the SSU has ample capacity to focus on the intertemporal externalities related to the capacity-constrained FSU. This is a reasonable assumption in many empirical settings, e.g., the ICU admission problem studied in Section 3. The system flow is depicted in Figure 1.

At the beginning of period t, the system state is given by a three dimensional vector,

$$s_t = (n^G_{l,t}, n^G_{h,t}, n^F_t),$$

where $n^G_{i,t}$ is the number of class i customers at the GK, $i \in \{l, h\}$, and n^F_t is the number of customers currently in the FSU. By the capacity constraints, $n^G_{i,t} \leq Q_i$ for $i \in \{l, h\}$ and $n^F_t \leq B$. For each customer

Figure 1 Overview of the System Flow

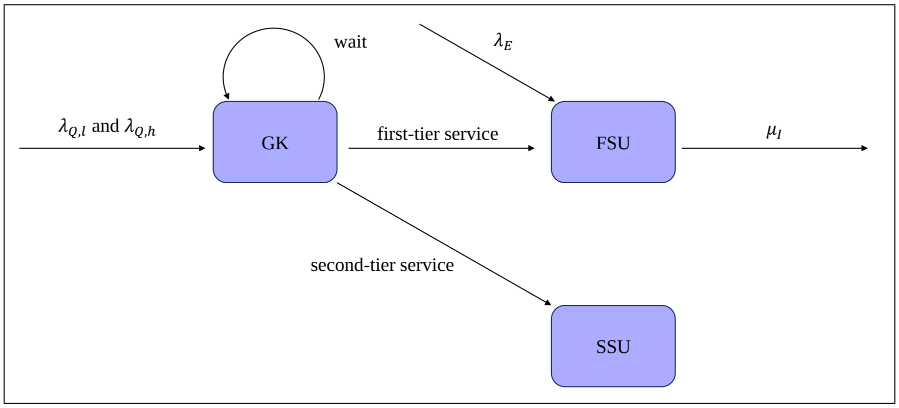

currently at the GK, the decision-maker chooses one of the following three destinations: the FSU, the SSU, or keeping them waiting at the GK for one more period. Since the customers are treated as identical within each class, the system's action can be described by the following four dimensional vector

$$d_t = (a_{l,t}, r_{l,t}, a_{h,t}, r_{h,t}),$$

where $a_{i,t}$ and $r_{i,t}$ denote the numbers of customers sent to the FSU and SSU of class $i \in \{l, h\}$, respectively. Due to the capacity constraint in the FSU, the admissible action set for system state s_t is

$$\Pi(s_t) = \left\{(a_{l,t}, r_{l,t}, a_{h,t}, r_{h,t}) : a_{l,t} + r_{l,t} \le n^G_{l,t},\ a_{h,t} + r_{h,t} \le n^G_{h,t},\ a_{l,t} + a_{h,t} \le B - n^F_t\right\}. \tag{1}$$

$\Pi(s_t)$ specifies the following set of constraints: The first two constraints state that the total number of customers sent to the FSU and SSU must be smaller than or equal to the total number of customers currently at the GK. The last constraint requires that the total number of customers sent to the FSU must be smaller than or equal to the current number of available servers in the FSU, i.e., $B - n^F_t$.

Let $u_{a,l}$ and $u_{a,h}$ be the system's expected utilities of sending a low and high-priority class customer to the FSU, respectively. In addition, denote the expected utilities of sending each low and high-priority customer to the SSU by $u_{r,l}$ and $u_{r,h}$, and the expected utilities of keeping each low and high-priority customer waiting at the GK by $u_{w,l}$ and $u_{w,h}$, respectively. These utility parameters incorporate all relevant factors that may affect the decision, such as customer demand, financial profit, and operational constraints. Such setting is standard in dynamic discrete choice models to estimate decision from observed data.

The system's expected utility $u(s_t, d_t)$ associated with state s_t and action d_t in period t is given by

$$u(s_t, d_t) = u_{a,l}a_{l,t} + u_{r,l}r_{l,t} + u_{w,l}\left(n^G_{l,t} - a_{l,t} - r_{l,t}\right) + u_{a,h}a_{h,t} + u_{r,h}r_{h,t} + u_{w,h}\left(n^G_{h,t} - a_{h,t} - r_{h,t}\right), \tag{2}$$

which is the sum of FSU routing, SSU routing, and waiting utilities for the two classes of customers. We assume that the system's total per period utility can be written as

$$U(s_t, d_t, \varepsilon_t) = u(s_t, d_t) + \varepsilon_t(d_t). \tag{3}$$

In the above, $\varepsilon_t(d_t)$ is the idiosyncratic utility component associated with action d_t that deviates from the average utility of the action for the two customer classes. For example, it can reflect the impact of customer-specific factors. This random utility shock is observed by the system when making decision, but not the researcher.

At the beginning of period t, the system observes the current state s_t and the idiosyncratic utility component ε_t, then it chooses the optimal action d_t that solves the following infinite horizon utility maximization problem:

$$\sup_{d_t \in \Pi(s_t)} \mathsf{E}\left\{\sum_{j=t}^{\infty} \beta^{j-t} U(s_j, d_j, \varepsilon_j) | s_t, \varepsilon_t\right\}. \tag{4}$$

The discount factor $\beta \in (0,1)$ captures the trade-off between the system's perceived utility of current customers and future customers, which is the focus of our study. It represents the relative weight of the next period utility in the objective. The expectation in (4) is taken over both the future random component ε_j and the transition of system state–i.e., the arrivals, departures, and routing of customers in each period. In addition, it is conditional on both s_t and ε_t, as the random component is observable to the system before making a decision in period t. In line with the structural estimation literature, the optimization problem models the system's decision making process under the observed flows of customers. Thus, the action utilities and discount factor should be perceived as behavioral parameters for estimating how the system balances complex and partially conflicting goals according to the data.

A key feature of our set-up is that the per-period utility parameters, $\{u_{a,\iota}, u_{r,\iota}, u_{w,\iota}\}$ for $\iota \in \{l, h\}$, are constant across states. In particular, the utility of admitting a customer to FSU is not affected by the current FSU occupancy n_t^F. This assumption is key to our identification—it implies that the impact of system state on routing decisions operates through the future continuation value instead of the current period payoff. Accordingly, the discount factor measures how the system balances current versus future utilities. We discuss the empirical plausibility of this assumption in our ICU setting in Section 4.2 and Section OS.3 of the Online Supplement.

After the system chooses an action in period t, the number of GK customers of class i becomes $n_{i,t}^G - a_{i,t} - r_{i,t}$, and the number of customers in the FSU is $n_t^F + a_{l,t} + a_{h,t}$. The system state then evolves according to the following two steps. First, $A_{i,t}$ new customers of class i arrive to the GK, and E_t customers arrive to the FSU through other channels–i.e., the external arrivals. If the GK or FSU is full, new arrivals can not be accepted. Thus, the total accepted GK and FSU arrivals are given by $A_{i,t}^{acc} = \min\left\{A_{i,t}, Q_i - (n_{i,t}^G - a_{i,t} - r_{i,t})\right\}$ for $i \in \{l, h\}$ and $E_t^{acc} = \min\left\{E_t, B - (n_t^F + a_{l,t} + a_{h,t})\right\}$, respectively. Second, D_t customers leave the FSU as they complete their first-tier service. This completes the system transition for period t. The system state at the beginning of period $t+1$ is thus given by $s_{t+1} = (n_{l,t+1}^G, n_{h,t+1}^G, n_{t+1}^F)$, with

$$n_{i,t+1}^G = n_{i,t}^G - a_{i,t} - r_{i,t} + A_{i,t}^{acc} \quad \text{and} \quad n_{t+1}^F = n_t^F + a_{l,t} + a_{h,t} + E_t^{acc} - D_t. \tag{5}$$

It is clear that, by construction, the transition of s_t is Markovian, and its distribution only depends on s_t and d_t, but not ε_t. The timeline of the system transition is summarized in Figure 2.

Figure 2 **Timeline of System Evolution: depiction of how the state evolves within a single time-slot.**

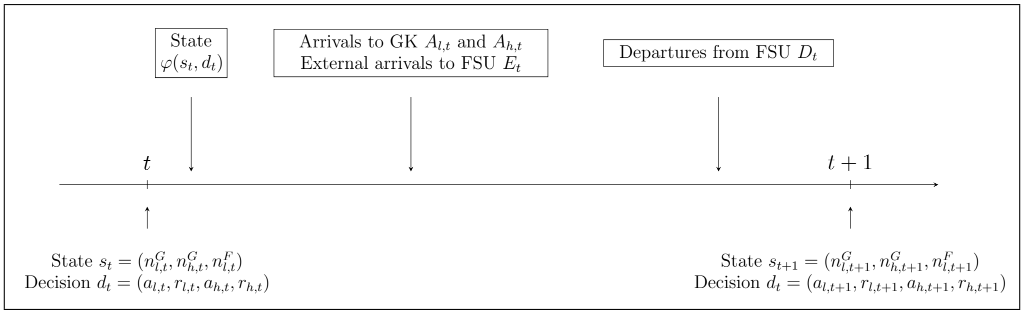

Both the current period's utility and how an action impacts the system state and future utilities may influence the system's decision in the current period. Define the value function as the objective in (4) given the optimal action sequence, i.e.,

$$V(s_t, \varepsilon_t) = \sup_{d_t \in \Pi(s_t)} \mathsf{E}\left\{ \sum_{j=t}^{\infty} \beta^{j-t} U(s_j, d_j, \varepsilon_j) | s_t, \varepsilon_t \right\}. \tag{6}$$

By (4), the system chooses the action that maximizes

$$d_t = \arg\max_{d \in \Pi(s_t)} (u(s_t, d) + \varepsilon_t(d) + \beta \mathsf{E}[V(s_{t+1}, \varepsilon_{t+1})]).$$

The last term on the right-hand side is the expectation of the future value function after the current action is taken. Thus, the optimal value function $V(s_t, \varepsilon_t)$ solves the following Bellman equation

$$V(s_t, \varepsilon_t) = \max_{d \in \Pi(s_t)} (u(s_t, d) + \varepsilon_t(d) + \beta \mathsf{E}[V(s_{t+1}, \varepsilon_{t+1})]),$$

where the expectation is taken over both the system transition to s_{t+1} and the random component ε_{t+1}.

The above Bellman equation is hard to evaluate due to the infinite state space associated with ε_t. Thus, we simplify the model by making the same conditional independence (CI) assumption as in Rust (1987).

Assumption 1 (CI) *The transition probabilities of the controlled process* (s_t, ε_t) *can be factored as*

$$\Pr(s_{t+1}, \varepsilon_{t+1} | s_t, \varepsilon_t, d_t) = q(\varepsilon_{t+1} | s_{t+1}) g(s_{t+1} | \varphi(s_t, d_t)), \tag{7}$$

where

$$\varphi(s_t, d_t) = \left(n^G_{l,t} - a_{l,t} - r_{l,t}, n^G_{h,t} - a_{h,t} - r_{h,t}, n^F_t + a_{l,t} + a_{h,t}\right) \tag{8}$$

denotes the "intermediary" state right after the action d_t is taken. The transition probability $g(s_{t+1} | \varphi(s_t, d_t))$ captures the random arrivals and departures shown by (5). Assumption (CI) states that s_{t+1} is sufficient to determine the distribution of ε_{t+1}. In other words, the random component $\{\varepsilon_t\}$ is superimposed on the state process $\{s_t\}$. In addition, following the literature, we assume that ε_t is independent and identically distributed (i.i.d.) and follows a type I extreme value distribution for each action $d \in \Pi(s_t)$. As shown in Rust (1987), this assumption leads to a closed-form expression of the conditional choice probability for action d_t given state s_t, as denoted by $f(d_t|s_t)$.

Proposition 1 *With the above set-up, the conditional choice probability for action d_t given state s_t has the following closed-form representation:*

$$f(d_t|s_t) = \frac{\exp\left(u(s_t,d_t) + \beta\tilde{V}(\varphi(s_t,d_t))\right)}{\sum_{d\in\Pi(s_t)}\exp\left(u(s_t,d) + \beta\tilde{V}(\varphi(s_t,d))\right)}, \tag{9}$$

where function $\varphi(s_t,d_t)$ is given by (8). The function $\tilde{V}(s)$ is the unique fixed point to the following functional equation

$$\tilde{V}(s) = \sum_{s'} \ln\left\{\sum_{d'\in\Pi(s')}\exp\left(u(s',d') + \beta\tilde{V}(\varphi(s',d'))\right)\right\} g(s'|s). \tag{10}$$

The explicit expression for $g(s'|s)$, i.e. the transition probability to state s' given s (the system state after the action is taken but before the random arrivals and departures take place), is provided in Section EC.1.1.

As in most dynamic discrete choice models, the choice probability (9) has a closed-form representation, and the value function $\tilde{V}(s)$ solves the functional equation (10). The new value function $\tilde{V}(s)$ represents the expected future utility given s, the intermediary state after the action has been taken. In the proposition below, we show that it is monotonically non-increasing in the number of customers in the FSU.

Proposition 2 *For two intermediary states (after actions are taken) s and s' with $(n_i^G) = (n_i^G)'$ for $i \in \{l,h\}$ and $(n^F) \le (n^F)'$, we have*

$$\tilde{V}(s) \ge \tilde{V}(s').$$

That is, for any given number of customers at the GK, the function $\tilde{V}(s)$ is monotonically non-increasing in the number of FSU customers.

PROOF: See Section OS.1.1. □

This proposition shows that, as the system routes more customers to the capacity-constrained FSU, the future expected utility decreases. Thus, when routing customers, the system needs to balance the impact on current customers as well as its capacity to serve future customers.

2.2. Identification of Discount Factor

In this study, we identify the utility parameters and the discount factor in the model *jointly* from data. It is well known that identifying the utility parameters and the discount factor jointly is in general very difficult (e.g. Lemma 3.3 in Rust 1994 or Proposition 2 in Magnac and Thesmar 2002). The reason is that multiple combinations of discount factor and utility parameters might lead to the same choice probabilities in (9) for all states and actions, i.e., they are observationally equivalent. Due to such challenge, most empirical studies assume the discount factor is known (e.g., 0.975) and estimate the utility parameters only. However, the pre-specified discount factors usually lack empirical support and economic justifications, as the implied discount rate can vary substantially across different settings (Frederick et al. 2002).

To identify the discount factor in our model, we leverage the theoretical development in Komarova et al. (2018). As highlighted in their paper, there is no prior identification results involving the discount factor for general parametric dynamic discrete choice models. The main steps for identification in Komarova et al. (2018) proceed as follows. They consider a dynamic discrete choice model where the utility function is linear in the unknown parameters. The choice set and transition probabilities are nonparametrically identified. For a given discount factor β, they first construct estimates for the utility parameters using the standard two-step estimation procedure pioneered by Hotz and Miller (1993). This reduces the identification problem to a one-dimensional search for $\beta \in (0, 1)$: If there is a unique value of β, together with the estimated utility parameters, that minimizes the objective function, the model (discount factor and utility parameters) can be empirically identified under some rank condition (see Theorem 1 in Komarova et al. 2018).

We extend and apply the identification results in Komarova et al. (2018) to our setting. By construction, our model satisfies the basic assumptions in Komarova et al. (2018), i.e., additive separability of utility, conditional independence of transition, and finite state space. By (2) and (3), the deterministic part of the per-period utility $u(s_t, d_t)$ is linear in the utility parameters $\{u_{a,\iota}, u_{w,\iota}, u_{r,\iota}\}$ for $\iota \in \{l, h\}$. Thus, the linear-in-parameter assumption is also satisfied. The details of these assumptions in our setting are discussed in Section EC.1.2. We construct a one-dimensional criterion similar to the one in Komarova et al. (2018) based on the maximum likelihood estimator in Rust (1987). We provide more details of the algorithmic approach in the next section.

An additional challenge for identifying the discount factor in our structural model is that our model has state-dependent action sets, as given by $\Pi(s_t)$. Thus, some actions may be infeasible in certain states. This differs from the setting in Komarova et al. (2018), which assumes the same admissible action set in all states. However, we note that for each state s_t, the admissible action set can be fully determined by (1). Thus, the identification results in Komarova et al. (2018) can be extended to our model, given that we use the admissible action set for each state according to (1). Accounting for the state-dependent action set is essential for our identification of discount factor. In our set-up, the intertemporal trade-off between current period and future payoffs stems from the fact that the FSU has limited capacity, which makes the action of admitting customer to FSU infeasible when the FSU gets full. If we assume the same action set in all states, such intertemporal trade-off vanishes in the model. This can lead to significant bias in the estimated discount factor. In Section EC.5, we use a simulation study to illustrate that allowing for state-dependent action sets is crucial to correctly identify and estimate the discount factor in our setting.

As pointed out in (Komarova et al. 2018), the key intuition for identification is that there is enough variation in the observed states and actions relative to the parameter space. This allows us to identify the parameters, including those in the utility function and the discount factor, from observed data. Thus, the linear utility assumption is essential because it limits the dimension of the parameter space. For instance, if we introduce a utility parameter for each state-action combination, the parameter space would be too large to achieve identification (e.g., Rust 1994, Magnac and Thesmar 2002). In Section EC.5, we develop a simple counterexample to show that when there is not sufficient variation in the observed states and actions, the discount factor cannot be identified.

The identification of the dynamic discrete choice model is relatively complex and does not allow explicit representations. That said, we include two examples below to provide some intuition of our identification strategy. In particular, we show that some aspects of the model can be directly identified by constructing appropriate *state-action pairs* and comparing their choice probabilities.

Proposition 3 *Assume the FSU admission utility is zero for both customer classes, i.e.,* $u_{a,h} = u_{a,l} = 0$.

(i) Consider two states $s_1 = (1, 0, 0)$ *and* $s_2 = (0, 1, 0)$. *The difference in SSU admission utilities can be identified as*

$$u_{r,l} - u_{r,h} = \ln\left(\frac{\Pr(r|s_1)}{\Pr(a|s_1)}\right) - \ln\left(\frac{\Pr(r|s_2)}{\Pr(a|s_2)}\right), \tag{11}$$

where $\Pr(a|s_1)$, $\Pr(r|s_2)$, $\Pr(a|s_1)$, *and* $\Pr(r|s_2)$ *denote the probabilities of admitting the customer to FSU and SSU under the two states, respectively.*

(ii) Consider two states $s_1 = (0, 1, 1)$ *and* $s_2 = (0, 1, 0)$. *Let* $\tilde{V}_k = \tilde{V}((0, 0, k))$ *denote the value function of the state with* k *patients in the FSU and no patients in the GK. Then we have*

$$\beta\big[(\tilde{V}_2 - \tilde{V}_1) - (\tilde{V}_1 - \tilde{V}_0)\big] = \ln\left(\frac{\Pr(a|s_1)}{\Pr(r|s_1)}\right) - \ln\left(\frac{\Pr(a|s_2)}{\Pr(r|s_2)}\right), \tag{12}$$

where $\Pr(a|s_1)$, $\Pr(r|s_2)$, $\Pr(a|s_1)$, *and* $\Pr(r|s_2)$ *denote the probabilities of admitting the high-priority customer to FSU and SSU under the two states, respectively.*

PROOF: See Section OS.1.2. □

The results can be interpreted as follows. By statement (i), the difference in the SSU admission utilities of the two customer classes can be identified using their ratios of SSU to FSU admission probabilities, i.e., $\ln[\Pr(r|s_i)/\Pr(a|s_i)]$. This is because the system is more likely to admit the customer to the SSU if the utility $u_{r,\iota}$ is high (given FSU admission utility is normalized to zero). For statement (ii), $\tilde{V}_2 - \tilde{V}_1$ (resp. $\tilde{V}_1 - \tilde{V}_0$) measures the impact on the future payoff from admitting one more customer to FSU when the FSU currently has one (resp. zero) customer. Thus, the left-hand side of (12) denotes the difference in such impacts when there is one customer in the FSU compared with when there is no customer in the FSU, multiplied by the discount factor β. Equation (12) shows that this difference can be identified by the change in the log ratios of FSU and SSU admission probabilities as the FSU occupancy increases from zero customer to one customer. Proposition 3 provides an example of the intuition behind our identification strategy that sufficient variation in observed states and actions leads to the joint identification of the utility parameters and the discount factor.

It is well known that to identify a dynamic discrete choice model, even with a given discount factor, the utility of one action needs to be normalized for all states (e.g., Rust 1994, Magnac and Thesmar 2002). The normalization level cannot be pinned down or refuted by the data without further restriction (e.g., Abbring and Daljord 2020). The common practice is to normalize the utility of a reference action to zero (see, e.g., Magnac and Thesmar 2002, Aguirregabiria and Mira 2010, Ryan 2012, Abbring and Daljord 2020). Depending on the model set-up and counterfactual changes, the normalization choice can affect the identification results and counterfactual estimates in general dynamic discrete choice models (see e.g., Magnac and Thesmar 2002, Norets and Tang 2014, Aguirregabiria and Suzuki 2014, Arcidiacono and Miller 2020). In our model, we normalize the utility of admitting a high-priority customer to the FSU by pre-specifying a value for $u_{a,h}$. The following proposition shows how the normalization choice impacts the identification of our structural model.

Proposition 4 *Assuming the GK has the same number of total customers $n^G_{l,t} + n^G_{h,t}$ at the beginning of each period, then the model identification, i.e., values of discount factor and differences in utility parameters, does not depend on the normalization level of FSU routing utility $u_{a,h}$.*

PROOF: See Section OS.1.3. □

Proposition 4 shows that when the number of total customers in GK is fixed in each period, the identified discount factor and utility differences are unaffected by the normalization choice of $u_{a,h}$. This extends the results in Komarova et al. (2018), which considers a constant action space in each period. In our general model set-up, the number of customers at GK can vary with time due to the routing decisions and random arrivals. Then, the model identification will generally hinge on the normalization choice. This is because the future states of GK, and thus future system utilities, will depend on the action taken in the current period (Abbring and Daljord 2020). As such, the normalization level affects the underlying trade-off between the current versus future customers. In these cases, the normalization choice should be justified based on the relevant objectives and constraints specific to the empirical setting (e.g., Ryan 2012).

2.3. Estimation Procedure

We now describe how our dynamic discrete choice model is estimated using the standard choice data, where we observe the system state and routing decisions of each customer. Suppose we have the state-action sequences (s_t, d_t) from the system for $t = 1, 2, ..., T$. First, the arrival and departure rates $\{\lambda_{Q,l}, \lambda_{Q,h}, \lambda_E, \mu_I\}$, as well as the FSU and GK capacities $\{Q_l, Q_h, B\}$ are estimated directly from data –

outside of the structural model. We provide more details about their estimation in our empirical setting in Section EC.2.2.

The other parameters — the discount factor, the utility of routing the low-priority customer to the FSU, and the utilities of routing customers to the SSU and keeping them waiting in the GK for the two classes — are estimated within the structural model. Denote these parameters by $\theta = \{\beta, u_{a,l}, u_{r,l}, u_{r,h}, u_{w,l}, u_{w,h}\}$. Given the observed states and actions, the likelihood for a fixed set of parameters, θ, is given by

$$l^f(\mathbf{s},\mathbf{d}|\theta) = \prod_{t=1}^{T} f(d_t|s_t,\theta) g(s_{t+1}|\varphi(s_t,d_t)) = \underbrace{\prod_{t=1}^{T} f(d_t|s_t,\theta)}_{l^d(\mathbf{s},\mathbf{d}|\theta)} \times \underbrace{\prod_{t=1}^{T} g(s_{t+1}|\varphi(s_t,d_t))}_{l^s(\mathbf{s},\mathbf{d})}, \tag{13}$$

where $(\mathbf{s},\mathbf{d})$ denotes the observed state and action sequences, i.e., $\{s_t, d_t\}$ for $t = 1, 2, ..., T$; $f(d_t|s_t,\theta)$ is the choice probability in (9) given parameter θ. The state transition probability $g(s_{t+1}|\varphi(s_t,d_t))$ is explicitly given in the Section EC.1.1. The likelihood l^f can be decomposed to two parts: $l^d(\mathbf{s},\mathbf{d}|\theta) := \prod_{t=1}^{T} f(d_t|s_t,\theta)$ is the part associated with choice probabilities, and $l^s(\mathbf{s},\mathbf{d}) := \prod_{t=1}^{T} g(s_{t+1}|\varphi(s_t,d_t))$ is the part from state transitions. We see that the structural parameter θ is only involved in the first part l^d, but not the second part l^s which only depends on the arrival and departure rates. Thus, we only need to maximize $l^d(\mathbf{s},\mathbf{d}|\theta)$.

We employ the nested fixed-point algorithm in Rust (1987) to estimate the utility parameters (conditioning on discount factor) by maximizing the likelihood of observed choices. The estimation procedure consists of two loops: The "inner" loop computes the function $\tilde{V}$ for a fixed θ using value iteration, and the "outer" loop searches for the value of θ that maximizes the log-likelihood $\ln l^d$ in (13) by a gradient descent algorithm. To summarize, for a candidate value of $\beta \in (0,1)$, we estimate the utility parameters $\{u_{a,l}, u_{r,l}, u_{r,h}, u_{w,l}, u_{w,h}\}$ that maximize the log-likelihood $\ln l^d$. Then, we choose the discount factor and its associated utility parameters that leads to the largest likelihood among all candidate β. Finally, we verify that the rank condition in Komarova et al. (2018) holds by simulation using the estimated parameters. To examine the proportion of variation explained by the estimated model, we compute the McFadden's pseudo R-squared as: $R^2_{McF} = 1 - \ln l^d(\hat{\theta}) / \ln l^{null}$. Here l^{null} is the "null" likelihood from a multinomial logistic regression model with only intercept terms, i.e., the action probabilities do not depend on system states.

3. Empirical Study: ICU Admission Decisions

In this section, we apply our structural model and the identification method to a suitable and important empirical setting in healthcare operations management. In particular, we study the ICU admission decisions for patients in the emergency department (ED) in a large hospital network.

3.1. Data and Clinical Setting

We utilize a large data set from 21 hospitals in a large hospital network in California, US. The data contains 312,306 hospitalizations over a period of two years before the COVID-19 pandemic. All patients are covered by the same insurance program as the hospital network is vertically integrated. The hospitals cover a large geographic area and intra-hospital transfers are quite rare ($< 3\%$). As such, we treat each hospital independently. Each observation in our data corresponds to a single hospitalization. In the data, we have patient level information such as age, gender, admitting hospital, diagnosis, and multiple severity scores. In addition, we observe the admission and discharge time for each unit a patient stayed in during the hospitalization, as well as the type of care the unit provides – i.e., ICU, transitional care unit (TCU), general medical or surgical ward, operating room (OR), or the postanesthesia care unit (PAR).

In this study, we focus on the routing decisions for the patients admitted to a medical service via the ED. This patient cohort comprises the largest proportion of admitted ICU patients. In addition, unlike surgical patients who usually have scheduled arrivals with fixed care protocols, there is much variation in the admission decisions for the ED medical patients (e.g., Kim et al. 2015). The final study cohort for the ED medical patients consists of 164,166 hospitalizations for 22 hospitals over a horizon of two years.[1] The detailed data selection process is described in Section EC.2.1. In this patient cohort, 19,683 (12.0%) are admitted to the ICU, and the remaining are admitted to a non-ICU unit. We find that the admitted cohort has higher average severity scores than the full cohort. The ED boarding time is similar for the two cohorts. While we focus on the admission decisions for the ED medical patients, we utilize the data from all hospitalizations to compute the external arrivals to ICU and the real-time occupancy level of the ICU in each hospital. Among all ICU admissions, around 36.2% are directly from our study cohort described above, while external arrivals account for 63.8%. Thus, it is crucial to include the external arrivals when calculating ICU occupancy.

For an ED patient, the ICU admission decision is generally made as follows. After a patient is stabilized in the ED, the ED physician provides an initial assessment about the patient condition. If the patient may require ICU admission, an intensivist will be called to the ED for a consultation. The ultimate decision regarding the patient's disposition requires the communication and coordination of many people including the ED physicians, the ICU intensivist, and hospital administrators, as well as various system-level factors such as nurse staffing availability, diversion policies, or alternative interventions available. In this study, we refer to this composite decision maker as the "hospital" and study its perceived behavior according to observed data. When making admission decisions, the hospitals need to balance the utilities of current ED patients and the impact on future system state. As the ICU provides the highest level of care, swift admission generally benefits the patients who need ICU care. However, given the limited capacity and high occupancy of ICU, this may restrict access to ICU care for future, perhaps more severe, patients. Thus, the admission

[1] For Hospital 21, we find that its ICU capacity experienced a substantial change during the sample period. Thus, we split it into two parts, i.e., before and after the capacity change, and treat them as two hospitals in the estimation.

decision introduces intertemporal externalities via the system's limited capacity to treat future patients. This fits into our general framework in Section 2.

Our goal is to estimate how hospitals balance the current ED patient versus future system capacity trade-off from data and to evaluate the impact of such balancing behavior on ICU congestion management. A priori, it is not clear how hospitals actually balance this trade-off, though there is empirical evidence that hospitals tend to reduce the likelihood of ICU admission when ICU is congested (e.g., Kim et al. 2015). In addition, due to the complex nature of the decision-making process for ICU admissions (e.g., multiple parties involved) and potential behavioral biases (e.g., Kim et al. 2020), hospitals are likely not fully aware of their current balancing behaviors. Moreover, without the estimation results, it is difficult for hospitals to evaluate the impact of such balancing behaviors on ICU congestion and other performance metrics. Thus, our study is valuable in providing both empirical evidence to understand the hospitals' balancing behaviors in terms of current patient and future system capacity and an empirical framework to evaluate system changes under a complex healthcare setting. This reveals new opportunities to improve hospitals' long term operational performance in managing patient flow and capacity.

3.2. Descriptive and Reduced-form Evidence

We first provide some descriptive evidence for the hospitals' balancing behavior under the intertemporal externalities. The hospital's balancing behavior can be reflected by how its ICU admission decisions vary with ICU occupancy levels. If a hospital primarily focuses on the current ED patients when making decisions, we would expect the ICU admission probability to be insensitive to ICU occupancy. On the other hand, if a hospital is relatively concerned with the system's ability to treat future patients, the ICU admission probability would drop when the ICU gets more congested, as the hospital needs to reserve the limited ICU capacity for future patients (e.g., by stabilizing less severe patients in ED or wards). In Figure OS.2, we show the ICU admission probability with respect to ICU occupancy level in three representative hospitals. The shaded area denotes the 95% level confidence interval. From the figure, we can see the ICU admission probability generally drops as ICU occupancy increases. Moreover, the decrease happens well before the ICU is full. This provides descriptive evidence for hospitals' balancing behaviors.

We further develop a multinomial logistic model to estimate the ICU admission decisions from data. The model controls for a variety of factors that may affect the admission decisions, including patient characteristics, system states, as well as hospital and time fixed effects. We find that a higher ICU occupancy is associated with lower ICU admission probability for all hospitals combined and most individual hospitals. That is, the hospitals tend to slow their ICU admission when the ICU is more congested. We also find the effects of ICU occupancy vary substantially across hospitals, suggesting significant heterogeneity in their behaviors. The set-up of the multinomial logistic model and the estimated coefficients are provided in Section EC.3 of the Electronic Companion. The results reflect that, at least to some extent, the hospitals take the impact on future system state into consideration when making admission decisions. In subsequent sections, we investigate such discounting behaviors using our dynamic discrete choice model.

We note that the decline in ICU admission probability is also consistent with an alternative mechanism in which the clinical benefit of ICU admission itself diminishes as the unit becomes more congested (e.g., due to resource and staff limitations). Our structural model attributes such a pattern to hospitals' balancing behaviors regarding current patient utility and future system capacity. Distinguishing the two mechanisms requires empirical evidence on whether treatment quality significantly varies with ICU occupancy. We discuss this in Section 4.2 and explain why we focus on the balancing behavior due to intertemporal externalities.

3.3. Dynamic Discete Choice Model for ICU Admission Decisions

In this section, we apply the general structural model in Section 2 to the hospital's ICU admission process. For each patient currently in the ED, the hospital can make three decisions: admit the patient to the ICU, admit the patient to non-ICU units (mainly wards), or keep the patient waiting in the ED. In light of the admission process, we view the ED as the gatekeeper (GK) in our model, which sends patients to ICU or non-ICU units. The ICU is seen as the first-tier service unit (FSU), as it provides the highest level of care for the patients but has limited capacity. The non-ICU units provide a lower level of care and usually have ample capacity.[2] Thus, they are regarded as the second-tier service unit (SSU) in the model. The transfers from non-ICU units to ICU are captured by external arrivals to the ICU in our model.

We define each period to be a two-hour time interval, which provides a reasonable amount of time for transferring the patient from one unit to the next after an admission request is issued.[3] We partition the ED patients into two classes by their severity (LAPS2) score. The low severity class is defined as patients with the severity score below the 85th percentile of the sample distribution, and the high severity class as those above the 85th percentile. That is, all medical patients admitted via the ED are included in our structural model as one of the two classes. The high severity patients on average are more likely to require ICU service. Thus, the two severity classes correspond to the high-priority and low-priority customers in our general set-up.

We assume the hospital's utility maximization problem is formulated as (4), with the per period utility given by (3). We have six utility parameters for the low and high severity classes: the *average* utilities of admitting the patient to ICU ($u_{a,l}$ and $u_{a,h}$), admitting the patient to non-ICU units ($u_{r,l}$ and $u_{r,h}$), and keeping the patient waiting in ED ($u_{w,l}$ and $u_{w,h}$). As discussed in Section 2.1, the utility parameters

[2] This is a reasonable assumption as the proportion of periods where the ward occupancy exceeds 95% of its capacity is less than 0.8% in our sample.

[3] In our sample, 94% of the patients wait in the ED only for one or two periods (shorter than four hours). Thus, we expect potential deterioration effect to be small. In addition, with each period being a two-hour interval, it is rare in the data to have multiple ED patients admitted to the ICU simultaneously. Thus, the allocation policy (e.g., FCFS) has little impact on the system dynamics.

represent the *average* utilities of each action for a patient class. On the other hand, the term $\varepsilon_t(d_t)$ captures the idiosyncratic shocks to the utility from patient-specific conditions. Due to the curse of dimensionality, it is prohibitive to let the action utility depend on patient-level characteristics. We discuss this issue in Section EC.5 and perform a robustness check using three patient classes. Our main findings remain very similar. The utilities and discount factor in the model are perceived as behavioral parameters for understanding the hospitals' actions from observed data. They incorporate all relevant factors that may affect the hospital's admission decisions, such as patients' clinical outcomes, operational costs, financial revenue, system-level constraints, etc. This is standard in structural model literature (see, e.g., Olivares et al. 2008, Wang et al. 2019, Rath and Rajaram 2022). We do not model the action utilities as explicit functions of the observable factors. This is infeasible in the dynamic discrete choice model given the curse of dimensionality. Moreover, the data requirements to be able to explicitly estimate each factor independently would be prohibitive.

We assume that keeping patients waiting in the ED brings lower average utility than admitting patients to ICU or non-ICU for both severity classes, i.e., $u_{w,\iota} \leq u_{a,\iota}, u_{r,\iota}$ for $i \in \{l, h\}$. This is a natural assumption for the following reasons. In our setting, the ED waiting times refers to the time difference between the time the order to admit is made and the time the patient is eventually admitted to an inpatient unit – often referred to as *ED Boarding*. The ED doctors and nurses are trained to quickly assess patients and stabilize them, not to provide ongoing care which these patients need to receive in inpatient units. Therefore, a longer ED boarding time is generally undesirable and associated with negative clinical outcomes (e.g., Singer et al. 2011). As shown in Table OS.17, most patients (94%) in our sample wait in the ED only for one or two periods. Thus, we expect ED waiting to bring lower utility than ICU and non-ICU admission, as it is only a temporary state in the patient routing process due to friction. Note that the parameters $u_{w,l}$ and $u_{w,h}$ represent the average utilities associated with ED waiting. Thus, the fact that ED waiting may benefit some patients in certain periods does not conflict with the assumption that ED waiting brings lower average utility than ICU and non-ICU admissions.

On the other hand, we allow the non-ICU admission utilities to be higher or lower than the ICU admission utilities. Patients who are not critically ill can receive sufficient care in the ward, so admitting them to the ward can bring the hospital higher utilities. This can be seen by the fact that more than half of the patients are eventually admitted to the non-ICU units (ward) instead of ICU for both classes (66% for high severity and 92% for low severity).

As discussed in Section 2.3, we need to normalize the utility of one action in order to identify the model. In our setting, we normalize the utility of ICU admission for a high severity patient to zero, i.e., $u_{a,h} = 0$. Then, in principle, the ICU admission utility of low severity class $u_{a,l}$ can be identified using the states with both low and high severity patients in ED. However, we find that such states are relatively infrequent in our sample (only 10%), meaning that we lack the variation to identify the difference in ICU admission utilities of the two severity classes. Thus, in our main specification, we additionally assume the ICU admission utility to be the same for the high and low severity classes, i.e., $u_{a,l} = u_{a,h} = 0$. This normalization choice does not mean ICU admission is costless. Instead, it implies that all other utility parameters are assessed *relative* to ICU admission (see, e.g., Magnac and Thesmar 2002).

We note that it is impossible to pin down the normalization level from data, and setting zero utility for normalized actions is a common assumption in the literature (e.g., Magnac and Thesmar 2002, Ryan 2012, Abbring and Daljord 2020). We further justify this assumption based on our empirical context. In our setting, assuming zero ICU admission utilities is equivalent to setting a reference point of utility across ED states: admitting all ED patients to the ICU leads to the same average per-period utility (zero) regardless of the ED state. We expect this to be a reasonable assumption for two reasons. First, admitting a patient to the ICU is the best the hospital can do for the patient's clinical outcome, which is the hospital's primary goal. This is especially true for a vertically integrated hospital with capitated payments as in our study. Thus, ignoring the impact on ICU capacity, admitting all ED patients to ICU should lead to the same *perceived* utility regardless of the number of ED patients. Second, ICUs have very high fixed and operating costs (e.g., specialized equipment and fixed nurse-to-bed ratios), thus the marginal utility of admitting one patient should be small. Finally, if a patient is eventually admitted to the ICU, it is likely because his latest risk condition implies significant risk, regardless of the pre-assigned severity class. Therefore, we set same ICU admission utility for the two severity classes.

In Section EC.7, we show that our main identification results are robust to alternative normalization choices. In addition, under current normalization choice, the estimated dynamic discrete choice model has good explanatory power for hospitals' decisions and fit the observed data well. These results suggest that it is reasonable to use the ICU admission action as the reference point across ED states. For completeness, we summarize the major assumptions in our model and empirical study, as well as the justifications and robustness checks for these assumptions in Section EC.8.

4. Main Empirical Results

This section presents our main empirical results on how hospitals balance the trade-off between the utilities of current patients and the impact on system state when making ICU admission decisions.

4.1. Estimation for ICU Admission Decisions

We estimate the structural model for ICU admission decisions in two ways. First, we estimate the model for all hospitals combined. That is, we estimate one set of parameters that maximizes the sum of log-likelihood from all hospitals. Second, we estimate the model for each hospital individually. We choose to estimate the model in both ways for the following reasons. In the combined estimation, we can obtain more reliable estimates using a much larger sample. However, it does not provide information on how the discounting behavior varies across hospitals. For this purpose, we also estimate the model for each hospital individually to reveal potential heterogeneity across hospitals.

To reduce the computational burden, we restrict the potential values of the discount factor to a discrete grid $\beta = \{0.1, 0.2, ..., 0.9\}$. While coarse, this discrete grid is granular enough to distinguish how the hospital balances the current patient versus future system capacity when making decisions. We conduct a robustness check with a finer grid for select hospitals and find almost same results. We use the bootstrap method to get the standard errors of the estimates. For each hospital, we resample its data for 500 times by randomly selecting from its state-action pairs $\{s_t, d_t\}_{t=1}^T$ with replacement.[4] Each resampled set has the same number of periods as the original sample. We then estimate the model parameters for each of the 500 resampled sets to obtain the standard errors.

As described in Section 2.3, the system parameters (e.g., arrivals, capacities) are estimated outside the structural model. We describe the details of their estimation for our ICU admission problem in Section EC.2.2. In Table OS.6, we report the numbers of days and hospitalizations in our sample, ED and ICU capacities, average ICU occupancy, arrival and departure rates, and the proportion of patients that are eventually admitted to the ICU for the two classes of each hospital.

We find that the 22 hospitals in our study have very different sizes, work loads, and admission behaviors. For example, the ICU capacity varies from 7 to 36 beds across hospitals. The ICU admission probabilities also vary substantially across hospitals even for the same severity class. In addition, the ICUs in the hospitals are generally congested. The average ICU occupancy in most hospitals is higher than 50%, and, in some hospitals, higher than 70%. Finally, the ICU admission probability for the high severity class is above 30% for most hospitals, which is usually three or four times larger than that for the low severity class (around 8%). This implies that the admission decisions are very different for the two severity classes, which is captured in our structural model. However, with a much larger population, the low severity class contributes more (58%) to the total ICU admissions of ED medical patients than the high severity class (42%). Thus, it is important to include both severity classes in our model.

4.2. Estimated Discount Factor and Action Utilities

In what follows, we provide the main estimation results of the structural model, i.e., the estimated discount factor and the utility parameters. Table 1 reports the estimation results for all hospitals combined. The McFadden's pseudo R^2 of the structural model is 0.14, which is comparable to that from a comprehensive multinomial logistic regression, as will be discussed in the next section. The estimated discount factor is $\hat{\beta} = 0.3$, and is statistically significantly different from adjacent levels. In addition, we find that the likelihood function decreases monotonically as we deviate away from the estimated $\hat{\beta} = 0.3$, which further validates the identification and estimation results (see Figure OS.4).

At first glance, the estimated $\hat{\beta}$ is quite surprising. In most empirical studies with dynamic choice models, the discount factor is assumed to be relatively large, e.g., 0.975 or 0.99. However, we see that the estimated

[4] The log-likelihood depends on the model parameters only via the choice probabilities, which can be calculated explicitly for a given state-action pair (s_t, d_t) by (9).

Table 1 Estimation Results of Structural Model: All Hospitals Combined (N = 154,140 Hospital-periods)

Discount factor	Low Severity		High Severity		
$\hat{\beta}$	$\hat{u}_{w,l}$	$\hat{u}_{r,l}$	$\hat{u}_{w,h}$	$\hat{u}_{r,h}$	R^2
0.3	-0.071	1.950	-0.932	0.671	0.14
(0.005)	(0.001)	(0.010)	(0.018)	(0.013)	

Note: Standard error is reported in parenthesis.

$\hat{\beta}$ is much smaller than these levels. In our model, the relatively small value of $\hat{\beta}$ implies that the hospitals are *not* very "forward-looking" when making ICU admission decisions. Given each period in our model is a two-hour interval, it suggests that the hospitals mainly focus on the current patients and barely consider the impact of their decisions on the system beyond the next six hours (after three 2-hour periods, $0.3^3 \approx 0.03$). Alternatively, $\hat{\beta} = 0.3$ means that the current period utility accounts for 70% in the decision-maker's objective ($1/\sum_{t=0}^{+\infty} \beta^t = 1 - \beta$). The above findings highlight the importance of identifying the discount factor using real data instead of assuming a pre-specified value. They complement the results in De Groote and Verboven (2019) and Ching and Osborne (2020), which show that consumers are more impatient than usually assumed, i.e., having smaller discount factors.

Recall that the utility parameters represent the average perceived utilities measured relative to the ICU admission decision, which is normalized to have zero utilities. By Table 1, the non-ICU admission utilities $\hat{u}_{r,l}$ and $\hat{u}_{r,h}$ are positive and significant for both the low and high severity classes (1.950 and 0.671). This suggests that it is, on average, desirable to admit a patient to the ward units. Indeed, only few (about 32%/8% of high/low severity) patients are admitted to the ICU as its care is expensive and (likely) unnecessary. On the other hand, the ED waiting utilities $\hat{u}_{r,l}$ and $\hat{u}_{r,h}$ are negative and significant for both classes (-0.071 and -0.932), which is inline with our assumption. Finally, both the waiting and non-ICU admission utilities are significantly lower for the high severity class than for the low severity class. Thus, it is on average more undesirable to keep the high severity patients temporarily waiting in the ED or to send them to the ward units, reflecting that the high severity patients are more likely to require ICU care. The large differences in the action utilities highlight the importance of differentiating the two severity classes in the model.

Next, we estimate the structural model for each hospital individually. To save space, we report the estimated discount factor and utility parameters, their standard errors estimated by bootstrapping, as well as the McFadden's R-squared for each hospital in Table EC.2 of the Electronic Companion. The results reveal substantial heterogeneity in the estimated discount factors across hospitals. In particular, 13 out of the 22 hospitals have relatively small estimated discount factors $\hat{\beta} \in \{0.1, 0.2, 0.3\}$, five have medium discount factors $\hat{\beta} \in \{0.4, 0.5, 0.6\}$, and the other four have relatively large discount factors $\hat{\beta} \in \{0.7, 0.8, 0.9\}$. The estimation results show that the hospitals behave very differently when making ICU admission decisions. The hospitals with small discount factors focus primarily on the current ED patients. The others with relatively large discount factors account more for the impact of their admission decisions on the system's ability to treat future patients. Due to the smaller sample size, the standard errors of the estimates are larger than those in Table 1 for the combined estimation. That said, we find the estimated $\hat{\beta}$'s are generally accurate for individual hospitals (see the note of Table EC.2). For most hospitals, the log likelihood decreases as we deviate from the estimated discount factor (see examples in Figure OS.5). Finally, we find the patterns of the utility estimates for individual hospitals are largely similar to those for all hospitals combined.

The heterogeneity in the discount factors reveals a novel aspect of the practice variation observed in medical literature (Corallo et al. 2014). Given the complex nature of the ICU admission and potential behavioral bias of participants, the hospitals may not be aware of the implicit discounting level in their admission decisions. From this aspect, our identification of discount factors is valuable as it reveals a potential lever for hospitals to manage their patient flow. In Section EC.9, we use simulations to show that assuming a "wrong" pre-specified discount factor in the model can lead to substantial deviation in the estimates of ICU congestion: if we assume a large (resp. small) discount factor and only estimate the utility parameters, we would underestimate (resp. overestimate) the occurrences of high ICU congestion periods. Such deviations become more significant when the capacity of the ICU becomes more constrained. Thus, correctly identifying the discount factor from data is crucial for accurately estimating the system performance and evaluating the impact of different system changes. This is also highlighted in De Groote and Verboven (2019) and Ching and Osborne (2020) in very different settings to ours.

In our model, we interpret the observed behaviors of hospitals—the drop in ICU admission probability as the ICU occupancy increases—as reflecting how they balance the current patient versus future system capacity in ICU admission decisions. This forward-looking consideration is captured by the discount factor estimated from the data. We acknowledge that the observed behaviors can also arise from alternative mechanisms. In particular, if the utility of ICU admission declines as the ICU gets congested (e.g., due to reduced care quality), the hospitals may also lower their ICU admission probabilities, as the additional benefit from ICU admission over non-ICU units diminishes for some patients.

To examine whether such a mechanism could drive our results, in Section OS.3, we conduct a reduced-form analysis to determine if, for our patient cohort and setting, whether ICU occupancy is associated with patient outcomes. Our analysis suggest that neither the ICU occupancy level at the time of a patient's admission nor the average ICU occupancy during a patient's ICU stay has significant impact on a patient's mortality, readmission, or hospital LOS. This is, perhaps because, ICUs typically have dedicated medical resources for each bed and maintain a fixed nurse-to-bed ratio. Thus, the quality of care is unlikely to deteriorate significantly for most patients when the ICU is congested. Additionally, for some hospitals in our sample (e.g. Hosp 11 in Figure OS.2), the decline in ICU admission probability starts as early as at 20% occupancy – well before it is likely to impact outcomes–, suggesting congestion-dependent utilities is unlikely to be the main driver of admission behavior.

While we do not find empirical evidence for the mechanism of congestion-dependent treatment quality, we emphasize that our estimation results in Section OS.3 apply only to the ED medical patients in our study cohort, and they should not be interpreted as causal. Multiple studies have found that hospital unit congestion can negatively affect outcomes due to speed-up or early discharge (Kc and Terwiesch 2009, Kc and Terwiesch 2012, Chrusch et al. 2009, Halpern and Pastores 2010). Fully ruling out alternative mechanisms (e.g., congestion-dependent utilities) would require a more thorough (*causal*) investigation of the impact of ICU congestion on patient outcomes. Relatedly, our identification and interpretation of the discount factor hinge on the pre-specified model assumptions (e.g., linear utility function, state-independent utility parameters), as is common in the identification of dynamic discrete choice models. Violations of these assumptions would also impact the interpretation of our results.

4.3. Heterogeneity in Discount Factors

We further investigate the heterogeneity in the estimated discount factors of the 22 hospitals. To do this, we compute the cross-hospital correlation between the estimated discount factor and a comprehensive set of hospital and patient characteristics. This includes ED and ICU sizes, arrival rates of ED patients and external patients to ICU, and ICU departure rate. We also consider multiple measures for ICU congestion, including average and median ICU occupancy, as well as the proportion of periods with fewer than one or two available ICU beds. For patient characteristics, we include the average age, the average LAPS2 score, the proportion of patients of the high severity class, and the proportion of admissions to the ICU from external arrivals. Finally, we include the variation in system states (ED and ICU occupancy) and the four estimated utility parameters. The full list of the characteristics are reported in Table OS.7. We do not run a multivariate regression by including multiple hospital characteristics simultaneously. This is because we only have 22 hospitals (data points) in our sample and many hospital characteristics are highly correlated.

In Table 2, we report the correlations between the estimated discount factors and hospital characteristics that are statistically significant at the 10% level or lower. The correlations for all other hospital characteristics not reported in Table 2 are not statistically significant. We find consistent evidence that the discount factor tends to be higher for hospitals with a more congested ICU, as shown by a lower ICU departure rate, higher average or median occupancy levels, and more periods with only one or two available ICU beds. This can be interpreted as hospitals with more congested ICUs can benefit more from being system-focused. Thus, they are more likely to save the bed for future patients when the ICU occupancy is high. In addition, the estimated discount factor is positively correlated with the average ED waiting time (EDWait_l and EDWait_h) and the utility of ED waiting for the low severity class ($u_{w,l}$). This can be explained as follows: When the ED waiting utility is high, it suggests that the hospital is able to provide better care within the ED. In such cases, the hospital can keep patients for longer in ED (with necessary monitoring and proactive actions to avoid patients deteriorating significantly) and save ICU beds for future patients. This leads to a larger discount factor in the structural model.

Table 2 Correlations between Estimated Discount Factor ($\hat{\beta}$) and Hospital Characteristics from Data

μ_I	AvgICUOccu	MedICUOccu	$\Pr(n_t^I \geq B-1)$	$\Pr(n_t^I \geq B-2)$	EDWait$_l$	EDWait$_h$	$u_{w.l}$
-0.428^*	0.445^*	0.518^{**}	0.552^{**}	0.632^{**}	0.545^{**}	0.581^{**}	0.493^*

Note: $^*p < 0.05$, $^{**}p < 0.01$ and $^{***}p < 0.001$.

Finally, we show that our structural estimation approach and reduced-form analysis provide consistent evidence for the heterogeneity in hospitals' discounting behaviors. We calculate the correlation between the estimated $\hat{\beta}$ from the dynamic discrete choice model and the coefficient γ_{ICU} of ICU occupancy for the ICU admission decision from the multinomial logit model described in Section 3.2. Note that γ_{ICU} is negative for most hospitals (see Section EC.3), suggesting the likelihood of ICU admission decreases as ICU becomes more congested. We find that the estimated $\hat{\beta}$ and γ_{ICU} are highly negatively correlated for the 22 hospitals (with a correlation of -0.768, significant at the 0.1% level). This means that the hospitals with larger identified discount factors are indeed more responsive to the ICU occupancy level in their admission decisions according to the multinomial logistic model. This provides additional support for our structural model as a valid approach to estimate how hospitals balance the current patient versus future system capacity in admission decisions.

4.4. Goodness of Fit and Robustness Checks

We briefly discuss the goodness of fit of the estimated structural model and several robustness checks for our main findings. The detailed results are discussed in Section EC.7 of the Electronic Companion and Sections OS.2, OS.4, and OS.5 of the Online Supplement.

We show that our estimated structural model has good explanatory power for the hospitals' decisions and fits the data well. First, the structural model leads to comparable McFadden's pseudo R^2 to the comprehensive multinomial logistic model for all hospitals combined and for most individual hospitals. In addition, the heterogeneity in hospitals' discount factors is consistent with the estimated effects of ICU occupancy from the reduced-form analysis. We also show, by simulation, that the key system statistics produced by the structural model are close to those observed in the data for each hospital. Finally, we estimate the structural model using the first half of our sample only and show it has good out-of-sample performance.

We conduct comprehensive robustness checks for our main empirical findings. As discussed in Section 3.3, we normalize the ICU admission utilities of both severity classes to zero in our main specification. We perform two robustness checks for the normalization choice. First, we estimate the model for all hospitals combined using the sample with fixed ED states. Similar to Proposition 4, we can show that the identification is not affected by the normalization choice of $(u_{a,l}, u_{a,h})$ in a model with constant ED state $(n_{l,t}^G, n_{h,t}^G)$. The discount factor from the joint estimation is still $\hat{\beta} = 0.3$. Second, we estimate the model for each hospital separately under alternative normalization choices of ICU admission utilities. We vary $u_{a,h}$ and $u_{a,l}$ between -0.3 to 0.3 and consider the cases with both $u_{a,l} = u_{a,h}$ and $u_{a,l} \neq u_{a,h}$. This range is chosen as we expect the marginal utility of ICU admission is small. Besides, it is generally comparable to the estimated utility differences of ED waiting and non-ICU admission relative to ICU admission (Table 1) as well as the magnitude of the idiosyncratic shock (with a standard deviation of 1.28) in our model. Our main results remain very robust: For individual hospitals, the estimated discount factor $\hat{\beta}$ remains unchanged or only changes to adjacent levels in all cases.

We estimate the model with stratified samples based on flu vs. non-flu seasons and day vs. night periods. The latter is similar to the robustness check done in Yu et al. (2016). We still find similar identification results to our main specification. To further address potential seasonality issues (e.g., day and night can have different ED arrivals), we develop a new dynamic discrete model by accounting for the time trend in the system transition. The new model includes the time indicator as a new state variable, and allows the arrival and departure rates to vary for the 12 two-hour intervals in a day. We find that the estimated discount factors from the new model are very similar to our main specification for all hospitals combined and most individual hospitals.

In our setting, we cannot model the action utility as functions of patient-level characteristics (e.g., severity) due to the curse of dimensionality. In our main specification, we use a two-class model with low and high severity patients. This captures the heterogeneity in patients' need for ICU care. As a robustness check, we also estimate our dynamic discrete choice model with three patient classes, including low, medium, and high severity patients. We find that the estimated discount factors in the three-class model remain very similar to the main specification for all hospitals combined and most individual hospitals. In addition, the estimated ED waiting and non-ICU admission utilities also exhibit similar patterns. This suggests that our main findings are robust under the three-class model, and our main setting with two severity classes can effectively capture the patient heterogeneity in the sample.

5. Counterfactual Studies

One of the main advantages of the structural estimation approach is its ability to conduct counterfactual studies. In particular, structural estimation aims to recover the underlying primitives of the decision-making process that are invariant to system or policy changes (e.g., Ryan 2012, Aguirregabiria and Mira 2010). This allows us to quantify the impact of system changes on key performance metrics that cannot be obtained by descriptive or regression analysis. Our counterfactual results reveal that understanding hospitals' discounting behaviors provide new opportunities to improve the system's long-term operational performance. Even we do not provide prescriptive solutions, our findings are managerially valuable in the complex and resource-constrained settings like ICU admissions.

We focus on metrics for high ICU congestion, which can have a substantial impact on clinical outcomes. (e.g., Kim et al. 2015). We first consider the probability of high ICU congestion, which is defined as the proportion of periods with only one or no empty bed(s) in the ICU:

$$\Pr(\text{HighCgstn}) = \frac{1}{T}\sum_{t=1}^{T} \mathbf{1}\left\{n_t^F \geq B-1\right\}. \tag{14}$$

For all hospitals in our sample, operating above the 95% occupancy level implies that there is at most one empty bed. We then estimate the number of ICU patients who spend their ICU stay under highly congested states in a year by

$$\text{Pats HighCgstn} = \underbrace{\Pr(\text{HighCgstn}) \times 365 \times 12}_{\text{total number of high congestion periods}} \times \mu_I \times (B-1), \tag{15}$$

where the inverse of the average LOS, μ_I, estimates the number of patients each bed can serve in a period. We multiply the product by $B-1$ as there is at most one bed available during high congestion periods.

In addition, we examine the probability that an external arrival patient will balk upon arrival because there is no ICU bed available. This is given by

$$\Pr(\text{Balk}|\text{External Arrival}) = \frac{1}{T}\sum_{t=1}^{T} \mathbf{1}\left\{n_t^F = B\right\}. \tag{16}$$

Using a similar argument as for (15), the number of external arrivals who balk in a year can be estimated by

$$\text{Pats Balk} = \underbrace{\Pr(\text{Balk}|\text{External Arrival}) \times 365 \times 12}_{\text{total number of full ICU periods}} \times \lambda_E, \tag{17}$$

where λ_E is the external arrival rate in each two-hour period.

In our counterfactual studies, we use the parameters estimated for each hospital individually, which are reported in Table EC.2 of the Electronic Companion. All counterfactual estimates are obtained using 500 simulation runs; each run has the same number of periods as in the data plus a three-month warm-up period. To save space, we report the results for six selected hospitals below. The results for other hospitals are qualitatively similar and provided in Table OS.8 of the Online Supplement. We keep the estimated model primitives the same when evaluating counterfactual changes, this is consistent with the literature (e.g., Mehta et al. 2017, De Groote and Verboven 2019, Wang et al. 2019).

In Section EC.6, we evaluate the counterfactual impact of adding one ICU bed while holding arrival and service rates fixed. We find that adding one ICU bed significantly reduces ICU congestion. Importantly, the benefit tends to be larger for hospitals with higher average ICU occupancy levels and for those with larger estimated discount factors. This highlights that the discount factor is an important perspective in hospitals' capacity planning decisions, capturing a distinct and complementary dimension of admission behavior beyond average system utilization.

5.1. Increasing Discount Factor

In this section, we study an alternative intervention of modifying the hospital's discounting behavior. As discussed earlier, hospitals typically behave in a present-focused way in the complex ICU admission decision process, since the discount factor from the joint estimation is only 0.3. We now consider the counterfactual of increasing the discount factors from the current estimated level $\hat{\beta}$ to 0.9, while keeping the utility parameters and ICU capacity unchanged. Practically, this means that the hospitals would account for the decision utilities and system states over the next one to two days ($0.9^{12} \approx 0.28$), rather than the next several hours. For some hospitals, the estimated $\hat{\beta}$ is already 0.9, so this counterfactual has no impact on their system performance.

We find that changing β does not affect the proportion of patients eventually admitted to ICU, which primarily depends on the non-ICU admission utilities and the workload of the system (arrival and departure rates). However, increasing the discount factor can significantly impact the high ICU congestion measures. The results are summarized in Table 3 for the subset of six hospitals. The "β to 0.9" columns report the reduction in ICU congestion from increasing β to 0.9. To facilitate comparison, we also show the reduction in ICU congestion from adding an ICU bed ("$B+1$" columns), as given in Table EC.3. We see that increasing the discount factor *alone* can have a substantial effect on reducing ICU congestion. For some hospitals, this translates to 40 fewer patients exposed to high ICU congestion in a year. Such benefits are comparable to that of adding one ICU bed, which is extremely costly to implement. This highlights the importance of understanding the balancing behaviors of hospitals in ICU admission decisions and reveals a potential approach for hospitals to reduce ICU congestion. We discuss practical methods for hospitals to increase the discounting level and the corresponding results in the "Heuris" columns in the next section.

Table 3 Counterfactual Estimates of Impact when β Increases from $\hat{\beta}$ to 0.9 (Select Hospitals)

Hosp	$\hat{\beta}$	ΔDays HighCgstn (in days)			ΔPats HighCgstn (in patients)			ΔPats Balk (in patients)			Increase in ED Time (in hours)	
		β to 0.9	$B+1$	Heuris	β to 0.9	$B+1$	Heuris	β to 0.9	$B+1$	Heuris	Low Class	High Class
1	0.3	4.19	6.27	3.38	34.94	52.26	28.20	4.68	6.73	4.27	2.05	0.63
2	0.5	6.04	6.78	6.44	47.03	52.78	50.12	7.53	8.43	8.73	1.84	0.63
5	0.1	6.79	18.74	5.24	24.67	68.04	19.02	4.55	8.84	3.11	1.75	0.44
8	0.1	5.58	10.90	5.45	33.21	64.88	32.41	4.69	7.92	4.62	1.36	0.76
9	0.5	4.27	8.60	6.37	38.99	78.61	58.22	6.42	13.57	10.56	1.55	0.49
14	0.3	1.58	2.36	1.73	17.43	25.90	19.07	3.26	5.26	3.54	1.26	0.46

While increasing β helps to reduce ICU congestion, it leads to longer ED waiting time as hospitals would reduce the likelihood of ICU admission as the unit gets more congested. This highlights the fundamental tradeoff between the patient level and system level considerations that hospitals are constantly faced with. We quantify such effects via simulation. The last two columns in Table 3 report the increase in average ED waiting time for the two classes of patients as β increases to 0.9. We find that a larger β indeed increases the average ED waiting time for both classes of patients. For example, as β increases from 0.3 to 0.9 in Hospital 1, the average ED waiting time increases by two hours and 36 minutes (0.6 hour) for low and high severity classes, respectively. However, the increase in ED waiting time is much more significant for the low severity class than for the high severity class. This suggests that when the discount factor increases, hospitals generally prefer to delay the admissions of low severity patients than the high severity ones, as the latter on average are more likely to require ICU care. Such behavior can partially offset the negative impact of longer ED waiting times, as the increase disproportionately impacts the less severe patients.

5.2. Practical Nudges for Increasing the Discount Factor

We discuss potential practical nudges for hospitals to increase their discounting level in ICU admissions, i.e., to account more for the system-level impact on future patients of decisions. As the ICU admission is a complex process involving multiple stakeholders, there may be multiple reasons a hospital is now more present-focused in making admission decisions. In the behavioral economics literature, studies have shown that present-focused behaviors can arise due to the challenges in forecasting future payoffs (Gabaix and Laibson 2017), complexity of valuing delayed payments (Enke et al. 2023, Imas et al. 2022), and framing effects related to the gain and loss (Faralla et al. 2017). These factors are relevant in the ICU admission setting as decision makers (e.g., ED physicians, ICU intensivists) often face complex situations with limited information and time. For example, in hospitals, accurate system state information may not be updated as frequently across units (Kim et al. 2020). Besides, conventional medical training has put more emphasis on the gains related to the timely care of current patients. For some hospitals, medical and operational constraints in the ED may also force them to act in a present-focused manner.

The proper intervention to achieve more system-focused behavior will depend highly on the root cause of the current, less system-focused behavior. Although our data does not allow us to tease out the root causes, we discuss potential interventions in below. First, the hospital can educate and train the people involved about the benefits of system-focused behaviors. The counterfactual studies in previous section can provide valuable evidence and insight for this. Second, hospitals can improve its information system make accurate system state available to people involved in admission decisions. Additionally, hospitals can develop and employ more effective models to reduce the uncertainty in future workload of ED and ICU (e.g., Ang et al. 2016). If other resource limitations drive the behaviors, the hospital may take proactive measures in the ED to stabilize the patients and/or address these staffing limitations. This can decrease the need for ICU and lead to better medical and economic outcomes (e.g., Weingart et al. 2013).

To further shed light on this, we propose a data-driven, easy-to-implement heuristic policy that mimics the policy from the structural model with $\beta = 0.9$. This provides a simple *nudge* that can achieve most of the benefits in reducing ICU congestion generated by a larger discount factor. In the heuristic policy, we assume the decisions are made independently for each patient in the ED. The decision probabilities are restricted to be independent of the ED state and can only change once with respect to the ICU state. We divide the ICU states to the low and high congestion regimes by a threshold T_{ICU} for ICU congestion level. The probabilities for ICU admission, non-ICU admission, and ED waiting for each patient in every period are denoted by

$$\{p^{(1)}_{I,\iota},\ p^{(1)}_{N,\iota}, p^{(1)}_{W,\iota}\} \text{ if } n^I < T_{ICU} \text{ and } \{p^{(2)}_{I,\iota},\ p^{(2)}_{N,\iota}, p^{(2)}_{W,\iota}\} \text{ if } n^I \geq T_{ICU},$$

with $\iota \in \{l, h\}$ for the low and high severity patient classes. The three action probabilities in each state sum up to one. [5] The constant probabilities $p^{(i)}_{I,\iota}$, $p^{(i)}_{N,\iota}$, and $p^{(i)}_{W,\iota}$, as well as the threshold level T_{ICU} are parameters to be designed such that the heuristic policy mimics the structural model with $\beta = 0.9$. In particular, we use the structural model with $\beta = 0.9$ to generate a large sample of state-action pairs for each hospital with 500 simulation runs. Then we estimate the parameters in the heuristic policy for each hospital to maximize the log-likelihood of simulated actions.

Table 4 provides the estimated heuristic policy for the six select hospitals, including the threshold T_{ICU} for separating the low ($n^I < T_{ICU}$) and high ICU regimes ($n^I \geq T_{ICU}$), as well as the the ICU and non-ICU admission probabilities ($p^{(i)}_{I,\iota}$ and $p^{(i)}_{N,\iota}$) in the two ICU regimes for each patient class. We observe that the hospitals change their action probabilities only when the ICU occupancy is relatively high, i.e., the threshold T_{ICU} is close to the capacity B. In the high ICU regime, the ICU (resp. non-ICU) admission probabilities for both types of patients mildly drop (resp. increase). Thus, a practical way to implement "increasing the discount factor" is for the hospital to be more responsive to system state and adjust its admission decisions when the ICU occupancy is high.

Table 4 Heuristic Policy for Select Hospitals

Hosp	B	T_{ICU}	Low severity patients				High severity patients			
			Low ICU: $n^I < T_{ICU}$		High ICU: $n^I \geq T_{ICU}$		Low ICU: $n^I < T_{ICU}$		High ICU: $n^I \geq T_{ICU}$	
			$p^{(1)}_{I,l}$	$p^{(1)}_{N,l}$	$p^{(2)}_{I,l}$	$p^{(2)}_{N,l}$	$p^{(1)}_{I,h}$	$p^{(1)}_{N,h}$	$p^{(2)}_{I,h}$	$p^{(2)}_{N,h}$
1	21	20	0.051	0.396	0.032	0.413	0.263	0.370	0.195	0.435
2	26	24	0.049	0.397	0.036	0.410	0.216	0.387	0.178	0.435
5	11	10	0.052	0.450	0.041	0.465	0.300	0.433	0.249	0.493
8	16	15	0.075	0.478	0.052	0.501	0.249	0.386	0.193	0.446
9	22	20	0.036	0.458	0.028	0.466	0.204	0.455	0.172	0.492
14	36	33	0.040	0.528	0.033	0.532	0.230	0.478	0.208	0.500

We then quantify the performance of the heuristic policy using simulation. They are reported in the "Heuris" columns in Table 3. By the "β to 0.9" and "Heuris" columns, we see that the heuristic policy

[5] When some of the actions are not permissible given the ICU state (e.g., due to full ICU), we normalize the probabilities of other permissible actions to have a sum of one. When estimating the heuristic policy, we drop the periods when ICU is full, as no ICU admissions are feasible in such states.

achieve very similar benefits in reducing ICU congestion compared with increasing β to 0.9. Additionally, we note that the overall ICU admission probabilities and ED waiting times for the two classes of patients, as well as the average occupancy level of ICU, are very similar under the heuristic policy and the structural model with $\beta = 0.9$. This suggests the hospitals could reduce ICU congestion safely by being more system-focused via a nudge that is relatively straightforward to implement in practice.

6. Conclusions

In many capacity-constrained service systems, it is important to balance the needs of current customers and the system's ability to serve future customers. We shed light on this aspect using a structural estimation approach. We develop a dynamic discrete choice model with multiple service levels and customer types, and measure the decision-maker's balancing behavior using the intertemporal discount factor. We show that the discount factor and utility parameters can be identified and estimated jointly from standard choice data, which contrasts with the common practice of pre-setting a discount factor in the literature. As is common in the literature on dynamic discrete choice models, our identification strategy relies on pre-specified model assumptions, such as the linear utility function and constant utility parameters. Thus, our results do not apply to settings where these assumptions do not hold.

We apply our model to an important problem in healthcare operations management: the ICU admission decisions of ED patients. Using a large US hospitalization data set, we find that the estimated discount factor is relatively small, far below the levels usually assumed in the literature. In addition, there is much heterogeneity in the balancing behaviors across hospitals. Thus, it is important to identify and estimate discount factor from data in different empirical settings. We then use counterfactual simulations to show that the level of discount factor can significantly affect the benefits of capacity intervention in the ICU. In addition, increasing the discounting level in admission decision can effectively reduce ICU congestion, although it leads to longer ED boarding time for low-severity patients. We also propose a simple heuristic policy to implement such intervention.

In this study, we consider a specific type of service system and use a dynamic discrete choice model to describe it. Future research may develop suitable structural models for other types of service systems and establish identification results of key behavioral parameters therein. In order to focus on the identification of discount factor, we have to make a number of simplifications in our model to ensure tractability. This inevitably ignores many realistic features of the system. For example, the customer arrival and departure rates may be heterogeneous and state-dependent. It would be interesting to investigate in future research whether these features can be incorporated into structural models. Moreover, there are other operational decisions that can be made in managing the system (e.g., the demand-driven discharge in ICU). As such, another potential direction is to study the joint impact of these decisions and the trade-offs they introduce.

Our empirical study on the ICU admission decisions has several limitations that may be explored in future research. First, the data we use have no direct information on the ICU admission process, as we can only observe the final admission units. Thus, there are many factors which are not explicitly included in the model nor do we have data to understand their impact on the decisions. Second, the data is limited to hospitals within one healthcare system that uses a capitated payment model. Thus, it is possible that other payment models may also impact how the hospitals internalize intertemporal externalities. For example, a Fee-For-Service system may drive hospitals to be even more focused on the current patients. In addition, a more comprehensive investigation of how ICU congestion affects patient outcomes is needed to entirely rule out alternative mechanisms for our main economic findings. These are deferred to future research.

Acknowledgments

We thank Marla Gardner and John Greene for their help in preparing the data, along with the staff in the Division of Research and hospitals in Kaiser Permanente Northern California for their time and invaluable contributions to this research.

References

Abbring, Jaap H, Øystein Daljord. 2020. Identifying the discount factor in dynamic discrete choice models. *Quantitative Economics* **11**(2) 471–501.

Abbring, Jaap H, Øystein Daljord, Fedor Iskhakov. 2025. Identifying present-biased discount functions in dynamic discrete choice models. *arXiv preprint arXiv:2507.07286* .

Aguirregabiria, Victor, Pedro Mira. 2010. Dynamic discrete choice structural models: A survey. *Journal of Econometrics* **156**(1) 38–67.

Aguirregabiria, Victor, Junichi Suzuki. 2014. Identification and counterfactuals in dynamic models of market entry and exit. *Quantitative Marketing and Economics* **12**(3) 267–304.

Akşin, Zeynep, Barış Ata, Seyed Morteza Emadi, Che-Lin Su. 2013. Structural estimation of callers' delay sensitivity in call centers. *Management Science* **59**(12) 2727–2746.

Ang, Erjie, Sara Kwasnick, Mohsen Bayati, Erica L Plambeck, Michael Aratow. 2016. Accurate emergency department wait time prediction. *Manufacturing & Service Operations Management* **18**(1) 141–156.

Arcidiacono, Peter, Robert A Miller. 2020. Identifying dynamic discrete choice models off short panels. *Journal of Econometrics* **215**(2) 473–485.

Batt, Robert J, Christian Terwiesch. 2016. Early task initiation and other load-adaptive mechanisms in the emergency department. *Management Science* **63**(11) 3531–3551.

Bray, Robert L, Yuliang Yao, Yongrui Duan, Jiazhen Huo. 2019. Ration gaming and the bullwhip effect. *Operations Research* **67**(2) 453–467.

Chan, Carri W, Vivek F Farias, Gabriel J Escobar. 2016. The impact of delays on service times in the intensive care unit. *Management Science* **63**(7) 2049–2072.

Ching, Andrew T, Matthew Osborne. 2020. Identification and estimation of forward-looking behavior: The case of consumer stockpiling. *Marketing Science* .

Chrusch, Carla A, Kendiss P Olafson, Patricia M McMillan, Daniel E Roberts, Perry R Gray. 2009. High occupancy increases the risk of early death or readmission after transfer from intensive care. *Critical care medicine* **37**(10) 2753–2758.

Corallo, Ashley N, Ruth Croxford, David C Goodman, Elisabeth L Bryan, Divya Srivastava, Therese A Stukel. 2014. A systematic review of medical practice variation in oecd countries. *Health Policy* **114**(1) 5–14.

De Groote, Olivier, Frank Verboven. 2019. Subsidies and time discounting in new technology adoption: Evidence from solar photovoltaic systems. *American Economic Review* **109**(6) 2137–72.

Dong, Jing, Elad Yom-Tov, Galit B Yom-Tov. 2018. The impact of delay announcements on hospital network coordination and waiting times. *Management Science* **65**(5) 1969–1994.

Dubé, Jean-Pierre, Vanessa Alwan, Xinyao Kong. 2024. The identification of the dynamic discrete-choice model with stockpiling. *Available at SSRN* .

Dubé, Jean-Pierre, Günter J Hitsch, Pranav Jindal. 2014. The joint identification of utility and discount functions from stated choice data: An application to durable goods adoption. *Quantitative Marketing and Economics* **12**(4) 331–377.

Edbrooke, David L, Cosetta Minelli, Gary H Mills, Gaetano Iapichino, Angelo Pezzi, Davide Corbella, Philip Jacobs, Anne Lippert, Joergen Wiis, Antonio Pesenti, et al. 2011. Implications of ICU triage decisions on patient mortality: a cost-effectiveness analysis. *Critical Care* **15**(1) R56.

Emadi, Seyed Morteza, Bradley R Staats. 2020. A structural estimation approach to study agent attrition. *Management Science* **66**(9) 4071–4095.

Enke, Benjamin, Thomas Graeber, Ryan Oprea. 2023. Complexity and hyperbolic discounting. *NBER Working Paper 31047* .

Faralla, Valeria, Marco Novarese, Antonella Ardizzone. 2017. Framing effects in intertemporal choice: A nudge experiment. *Journal of Behavioral and Experimental Economics* **71** 13–25.

Frederick, Shane, George Loewenstein, Ted O'donoghue. 2002. Time discounting and time preference: A critical review. *Journal of Economic Literature* **40**(2) 351–401.

Freeman, Michael, Susan Robinson, Stefan Scholtes. 2021. Gatekeeping, fast and slow: An empirical study of referral errors in the emergency department. *Management Science* **67**(7) 4209–4232.

Freeman, Michael, Nicos Savva, Stefan Scholtes. 2016. Gatekeepers at work: An empirical analysis of a maternity unit. *Management Science* **63**(10) 3147–3167.

Gabaix, Xavier, David Laibson. 2017. Myopia and discounting. *NBER Working Paper* (w23254).

Halpern, Neil A, Stephen M Pastores. 2010. Critical care medicine in the united states 2000–2005: an analysis of bed numbers, occupancy rates, payer mix, and costs. *Critical Care Medicine* **38**(1) 65–71.

Halpern, Neil A, Stephen M Pastores. 2015. Critical care medicine beds, use, occupancy and costs in the united states: a methodological review. *Critical Care Medicine* **43**(11) 2452.

Hao, Yu, Hiroyuki Kasahara, Katsumi Shimotsu. 2025. Semiparametric identification of the discount factor and payoff function in dynamic discrete choice models. *arXiv preprint arXiv:2507.19814* .

Hathaway, Brett A, Evgeny Kagan, Maqbool Dada. 2023. The gatekeeper's dilemma:"when should I transfer this customer?". *Operations Research* **71**(3) 843–859.

Hotz, V Joseph, Robert A Miller. 1993. Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies* **60**(3) 497–529.

Imas, Alex, Michael A Kuhn, Vera Mironova. 2022. Waiting to choose: The role of deliberation in intertemporal choice. *American Economic Journal: Microeconomics* **14**(3) 414–440.

Kc, Diwas S, Christian Terwiesch. 2009. Impact of workload on service time and patient safety: An econometric analysis of hospital operations. *Management Science* **55**(9) 1486–1498.

Kc, Diwas Singh, Christian Terwiesch. 2012. An econometric analysis of patient flows in the cardiac intensive care unit. *Manufacturing & Service Operations Management* **14**(1) 50–65.

Kim, Song-Hee, Carri W Chan, Marcelo Olivares, Gabriel Escobar. 2015. ICU admission control: An empirical study of capacity allocation and its implication for patient outcomes. *Management Science* **61**(1) 19–38.

Kim, Song-Hee, Jordan Tong, Carol Peden. 2020. Admission control biases in hospital unit capacity management: How occupancy information hurdles and decision noise impact utilization. *Management Science* **66**(11) 5151–5170.

Komarova, Tatiana, Fabio Sanches, Daniel Silva Junior, Sorawoot Srisuma. 2018. Joint analysis of the discount factor and payoff parameters in dynamic discrete choice models. *Quantitative Economics* **9**(3) 1153–1194.

Li, Jun, Nelson Granados, Serguei Netessine. 2014. Are consumers strategic? structural estimation from the air-travel industry. *Management Science* **60**(9) 2114–2137.

Magnac, Thierry, David Thesmar. 2002. Identifying dynamic discrete decision processes. *Econometrica* **70**(2) 801–816.

Mehta, Nitin, Jian Ni, Kannan Srinivasan, Baohong Sun. 2017. A dynamic model of health insurance choices and healthcare consumption decisions. *Marketing Science* **36**(3) 338–360.

Norets, Andriy, Xun Tang. 2014. Semiparametric inference in dynamic binary choice models. *Review of Economic Studies* **81**(3) 1229–1262.

Olivares, Marcelo, Christian Terwiesch, Lydia Cassorla. 2008. Structural estimation of the newsvendor model: an application to reserving operating room time. *Management Science* **54**(1) 41–55.

Rath, Sandeep, Kumar Rajaram. 2022. Staff planning for hospitals with implicit cost estimation and stochastic optimization. *Production and Operations Management* **31**(3) 1271–1289.

Rust, John. 1987. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society* 999–1033.

Rust, John. 1994. Structural estimation of markov decision processes. *Handbook of Econometrics* **4** 3081–3143.

Ryan, Stephen P. 2012. The costs of environmental regulation in a concentrated industry. *Econometrica* **80**(3) 1019–1061.

Singer, Adam J, Henry C Thode Jr, Peter Viccellio, Jesse M Pines. 2011. The association between length of emergency department boarding and mortality. *Academic Emergency Medicine* **18**(12) 1324–1329.

Song, Hummy, Anita L Tucker, Ryan Graue, Sarah Moravick, Julius J Yang. 2020. Capacity pooling in hospitals: The hidden consequences of off-service placement. *Management Science* **66**(9) 3825–3842.

Wang, Guihua, Jun Li, Wallace J. Hopp, Franco L. Fazzalari, Steven F. Bolling. 2019. Using patient-specific quality information to unlock hidden healthcare capabilities. *Manufacturing & Service Operations Management* **21**(3) 582–601.

Weingart, Scott D, Robert L Sherwin, Lillian L Emlet, Isaac Tawil, Julie Mayglothling, Jon C Rittenberger. 2013. ED intensivists and ED intensive care units. *American Journal of Emergency Medicine* **31**(3) 617–620.

Yu, Qiuping, Gad Allon, Achal Bassamboo. 2016. How do delay announcements shape customer behavior? an empirical study. *Management Science* **63**(1) 1–20.

Zhang, Zhe George, Hsing Paul Luh, Chia-Hung Wang. 2011. Modeling security-check queues. *Management Science* **57**(11) 1979–1995.

Electronic Companion to "Structural Estimation of Intertemporal Externalities with Application in ICU Admissions"

Yiwen Shen
Department of ISOM, HKUST, yiwenshen@ust.hk

Carri W. Chan
Decision, Risk, and Operations, Columbia Business School, cwchan@columbia.edu

Fanyin Zheng
Decision, Risk, and Operations, Columbia Business School, fanyin.zheng@columbia.edu

Gabriel Escobar
Division of Research, Kaiser Permanente, gabriel.escobar@kp.org

---

EC.1. Formulae and Modeling Assumptions

EC.1.1. Explicit expressions for state transition probability $g(s'|s)$

We provide explicit expressions for the function $g(s'|s)$ used in Proposition 1, which is the transition probability from the intermediate state $s = (\tilde{n}^G_{l,t}, \tilde{n}^G_{h,t}, \tilde{n}^F_t)$ (after action is taken) to the state $s' = (n^G_{l,t+1}, n^G_{h,t+1}, n^F_{t+1})$ at the start of the next period (before action). As the GK arrivals are independent of the FSU external arrivals and departures, we have

$$g(s'|s) = g_{Q,l}\left(n^G_{l,t+1}|\tilde{n}^G_{l,t}\right) g_{Q,h}\left(n^G_{h,t+1}|\tilde{n}^G_{h,t}\right) g_F\left(n^F_{t+1}|\tilde{n}^F_t\right), \tag{EC.1}$$

where $g_{Q,l}$ and $g_{Q,h}$ denote the transition probabilities for the numbers of GK customers from low and high classes respectively, and g_F denotes the transition probability for the number of FSU customers.

For the GK transition probabilities $g_{Q,l}$ and $g_{Q,h}$, we only need to consider the new arrivals for class $i \in \{l, h\}$, which follow truncated Poisson distributions with rate $\lambda_{Q,i}$ and truncation by $M_{A,i}$ from above. Additionally accounting for the GK capacity constraint, the number of GK arrivals is capped by $\bar{A}_i = \min\{M_{A,i}, Q_i - \tilde{n}^G_{i,t}\}$. Thus, the transition probability can be computed as

$$g_{Q,i}(m|n) = \begin{cases} (\lambda_{Q,i})^{m-n} \exp(-\lambda_{Q,i})/(m-n)! & \text{if } n \le m < n + \min\{M_{A,i}, Q_i - n\} \\ \sum_{j=\bar{A}_i}^{+\infty} (\lambda_{Q,i})^j \exp(-\lambda_{Q,i})/j! & \text{if } m = n + \min\{M_{A,i}, Q_i - n\} \end{cases} \tag{EC.2}$$

and $g_{Q,i}(m|n) = 0$ elsewhere.

For the FSU transition probability g_F, we need to consider both external arrivals and departures. The number of external arrivals E_t follows a Poisson distribution with rate λ_E, and is capped by the remaining

FSU capacity $B-\tilde{n}_t^F$. With E_t external arrivals in the period, the FSU would have total $\tilde{n}_t^F+E_t$ customers. Then, the number of departures, D_t, follows a Binomial-$(\tilde{n}_t^F+E_t,\mu_I)$ distribution. Given the number of external arrivals, E_t, the number of departures follows by $D_t=\tilde{n}_t^F+E_t-n_{t+1}^F$. We note that E_t can be greater than $\max\{n_{t+1}^F-\tilde{n}_t^F,0\}$. Summing up the probability of all possible choices of the Poisson-distributed E_t (and Binomial-distributed D_t accordingly), the transition probability $g_F(m|n)$ is given by

$$\begin{aligned} g_F(m|n) = &\sum_{j=\max\{m-n,0\}}^{B-n-1} \lambda_E^j \frac{\exp(-\lambda_E)}{j!} \frac{(n+j)!}{(n+j-m)!m!} \mu_I^{n+j-m}(1-\mu_I)^m \\ &+ \left(\sum_{j=B-n}^{+\infty} \lambda_E^k \frac{\exp(-\lambda_E)}{j!}\right) \frac{B!}{(B-m)!m!} \mu_I^{B-m}(1-\mu_I)^m, \text{ for } 0\le m\le B. \end{aligned} \tag{EC.3}$$

Using (EC.2) and (EC.3), we obtain the explicit expression for state transition probability $g(s'|s)$ by (EC.1).

EC.1.2. Assumptions for Identification

This section documents the assumptions in Komarova et al. (2018), which are also satisfied by our model. In their paper, x and a denote the system state and action, respectively; ε is the random perturbation in utility.

Assumption 1 *(i) (Additive Separability) For all a, x, ε, the per-period utility follows:*

$$u(a,x,\varepsilon)=\pi(a,x)+\varepsilon(a).$$

(ii) (Conditional Independence) The transition distribution of the states has the following factorization for all x', ε', x, ε, a:

$$P(x',\varepsilon'|x,\varepsilon,a)=Q(\varepsilon')G(x'|x,a),$$

where Q is the cumulative distribution function of ε and G is the transition law of x_{t+1} conditioning on x_t and a_t. Furthermore, ε_t has finite first moments, and a positive, continuous, and bounded density.

(iii) (Finite Observed State) $X=\{1,\ldots,K\}$.

Assumption 2 (Linear-in-Parameter) *For all a, x:*

$$\pi(a,x;\theta)=\pi_0(a,x)+\theta^\top\pi_1(a,x),$$

where π_0 is a known real value function, π_1 is a known p-dimensional vector value function and θ is the p-dimensional unknown parameter.

In our setting, we have action utility given by (2) as

$$u(s_t,d_t)=u_{a,l}a_{l,t}+u_{r,l}r_{l,t}+u_{w,l}\left(n_{l,t}^G-a_{l,t}-r_{l,t}\right)+u_{a,h}a_{h,t}+u_{r,h}r_{h,t}+u_{w,h}\left(n_{h,t}^G-a_{h,t}-r_{h,t}\right).$$

This is the deterministic part of the per-period utility. Thus, it is indeed linear in parameters $\theta=\{u_{a,l},u_{r,l},u_{w,l},u_{a,h},u_{r,h},u_{w,h}\}$. In our model, the functions π_0 and π_1 specify to $\pi_0(d_t,s_t)\equiv 0$ and

$$\pi_1(d_t,s_t)=[a_{l,t},r_{l,t},n_{l,t}^G-a_{l,t}-r_{l,t},a_{h,t},r_{h,t},n_{h,t}^G-a_{h,t}-r_{h,t}]^\top.$$

EC.2. Data Selection and Estimation of System Parameters for ICU Admission

EC.2.1. Data Selection Process

We start from a total of 312,306 hospitalizations over two years. We restrict our study to the hospitalizations admitted to a medical service via the ED, which comprises the largest proportion of admitted patients. For the ED patients, they appear in our data set as soon as the admission decision has been made; as such, we do not have information about patients discharged home from the ED nor patients for whom a disposition decision has not yet been made. We drop 12 hospitalizations with unknown gender and 9,128 (4.8%) hospitalizations for patients who experience hospital transfers or transports outside of the hospital network. Our study focuses on three possible decisions for each patient in each decision epoch: keep the patient waiting in the ED, admit the patient to the ICU, or admit the patient to a non-ICU unit (e.g. the ward or TCU). We drop 3,066 (1.7%) hospitalizations where the patient was admitted to other units – e.g., OR or PAR, from the ED. Finally, we drop 1,675 (1%) hospitalizations with ED waiting time longer than 12 hours as these episodes can be considered outliers (the average waiting time is shorter than two hours).

We restrict our study cohort to the periods of each hospital with stable ICU capacity and occupancy. First, we discard the first and last month of data for all hospitals. Second, for several hospitals, we drop the period at either end of the sample where the ICU occupancy dramatically fluctuates or significantly differs from the more stable period in the middle. Finally, for hospital 21, we find that its ICU capacity experienced a substantial increase during the sample period (from 13 to 16). As a result, we split it into two parts, i.e., before and after the capacity change, and treat them as two hospitals in the estimation. We refer to 22 hospitals in our study cohort. The patient characteristics of the final study cohort and the ICU admission subset are summarized in Table OS.5 of the Online Supplement. The number of days and hospitalizations for each hospital in the final study cohort are summarized in Table OS.6 of the Online Supplement. In total, we drop 11,268 (6.4%) hospitalizations that are outside the stable periods.

EC.2.2. Estimation of System Parameters

As described in Section 2.3, the arrival and departure rates, as well as the unit capacities, are estimated directly from data – outside of the structural model. We now describe how we do this for the ICU admission problem. Recall we define each period as a two hour interval in the model. We estimate the ED arrival rates $\lambda_{Q,i}$ and maximum arrival number $M_{A,i}$ for $i \in \{l, h\}$ using the average and maximum number of arrivals to the ED for the two classes in each period. We estimate the ICU external arrival rate λ_E using the average number of patients admitted to ICU in each time slot who are not included in our low and high severity ED groups. The departure probability μ_I is estimated as the ratio of total number of departures to the total periods of ICU stay across all ICU patients. The ICU capacity B is set to be the maximum number of *all* patients (medical and surgical, emergency and elective) in the ICU observed from data. This is a reasonable assumption as ICU often operates under high congestion. We discuss the validity of these estimates in Section EC.8.

Our data captures the number of patients admitted to the hospital from the ED, but does not include any patients who are discharged from the ED. Thus, it is difficult to accurately determine the ED capacity in our model. However, we note that the number of ED boarding patients is relatively low (e.g., average 1.21 patients). Thus, we set ED capacities Q_i using the following heuristic: $Q_i = M_{Q,i} + \lfloor \sqrt{M_{A,i}} \rfloor$, where $M_{Q,i}$ is the maximum number of ED patients observed in the data; $M_{A,i}$ is the maximum number of arrivals in each period; and $\lfloor \cdot \rfloor$ denotes the floor function.The square root term $\sqrt{M_{A,i}}$ is introduced as a heuristic "safety buffer" to ensure we have ample ED capacity to avoid balking upon arrival to the ED, which rarely happens in reality. We verify by simulation that the ED rarely reaches its full capacity Q_i in our structural model. In addition, the choice probabilities are very robust to alternative choices of the ED capacity. The estimated system statistics for each hospital are reported in Table OS.6 of Online Supplement.

EC.3. Reduced-form Evidence for Discounting Behavior in ICU Admissions

In this section, we conduct reduced-form regressions to analyze the main determinants of the system's ICU admission decisions. This provides complementary evidence for the discounting behaviors of hospitals and serves as a benchmark for evaluating the explanatory power of our structural model. We apply a multinomial logistic model to estimate the ICU admission decisions. In each two-hour period, the hospital chooses one of the three options for each patient: admit to the ICU; admit to non-ICU units; or make the patient wait in the ED. At the start of each period, we construct system "snapshots" which include detailed information for each patient in the ED as well as the state of the ED and ICU. We include patient characteristics, system state variables, and seasonality effects as potential determinants of the admission decisions.

We use the non-ICU admission decision as the base case, and estimate the probabilities of the ICU admission ($d = \text{ICUAdm}$) and waiting ($d = \text{Wait}$) decisions relative to the non-ICU admission decision respectively. For patient i who is in the ED at the start of period t, we set up the multinomial logistic model:

$$\ln\left[\frac{\Pr(d_{it}|\mathbf{X}_i, \mathbf{S}_t)}{\Pr(\text{nonICU}_{it}|\mathbf{X}_i, \mathbf{S}_t)}\right] = \gamma_{0,d} + \gamma_{L,d}\text{LAPS2}_i + \gamma_{ICU,d}\text{ICUOccu}_t + \gamma_{ED,d}\text{EDNum}_t + \boldsymbol{\gamma}_{Z,d}^{\top}\mathbf{Z}_{i,t} + \epsilon_{it}, \tag{EC.4}$$

where $\Pr(d_{it})$ and $\Pr(\text{nonICU}_{it}|\mathbf{X}_i\mathbf{S}_t)$ denote the conditional probability of action d_{it} and admitting patient i to non-ICU units in period t, respectively. In (EC.4), LAPS2 is the patient's main severity score. ICUOccu_t denotes the current ICU occupancy level. As the ICU sizes vary dramatically across the hospitals, we use the ICU percentile rank to measure occupancy. EDNum_t denotes the number of current ED patients. $\mathbf{Z}_{i,t}$ denotes other control covariates, given by

$$\mathbf{Z}_{i,t} = \{\text{Gender}_i, \text{Age}_i, \text{COPS2}_i, \text{CHMR}_i, \text{Hosp}_i, \text{DepPre}_t, \text{AvgLAPS2}_t, \text{DayOfWeek}_t, \text{HourOfDay}_t, \text{Month}_t\}.$$

It includes patient i's gender, age, identifier for hospital admitted, as well as other two severity scores COPS2 and CHMR. In addition, it includes several system state variables: the number of patients who left the ICU in the previous period (DepPre_t), the average severity level measured by the LAPS2 score of the current ICU patients in period t (AvgLAPS2_t), and several categorical variables to capture the potential seasonality and time trend in the decisions. Here DayOfWeek_t and HourOfDay_t denote the day of week and hour of day respectively; Month_t is the dummy variable representing the month in the sample (total 23 months). To account for the heteroskedasticity, we cluster standard errors by hospitals in the regression.

We first estimate the model by combining the patient data from all hospitals. We use the McFadden's pseudo R-squared to measure the goodness of fit of the model. It is defined as $R^2 = 1 - \ln l^{mod}/\ln l^{null}$, where l^{mod} is the likelihood from the estimated model and l^{null} is the likelihood from the "null" model that only includes the intercept and a categorical variable for each hospital in (EC.4).

Table EC.1 Estimation results for Multinomial-Logistic Regression (EC.4), N = 183,691, R-squared = 0.16

	LAPS2_i	ICUOccu_t	EDNum_t
	γ_L	γ_{ICU}	γ_{ED}
Waiting	0.008	1.178	0.241
	(0.000)	(0.031)	(0.006)
ICUAdm	0.028	-0.396	-0.014
	(0.000)	(0.030)	(0.006)

Standard error is reported in parenthesis.

Table EC.1 reports the estimated coefficients for three main variables in the multinomial logistic model: LAPS2_i, ICUOccu_t, and EDNum_t. All the coefficients are statistically significant and have the expected sign. In particular, higher severity score (LAPS2) increases the probability of admission to ICU relative to other units. In addition, the estimates of γ_{ICU} and γ_{ED} suggest that, even after controlling for patient characteristics and fixed effects, a busier system state (more congested ICU or ED) decreases the probability of ICU admission and increases the probability of waiting. Such evidence suggests that the hospitals indeed internalizes the intertemporal externalities on the ICU admission decisions by adjusting their behaviors according to the current system state. The McFadden's pseudo R^2 for the combined estimation is 0.16, consistent with the magnitude seen for models of operational decisions in healthcare systems.

Considering the heterogeneity across hospitals, we further estimate model (EC.4) for each hospital separately after dropping the hospital dummy Hosp_i. As a robustness check, we also estimate a multinomial logistic model with only three variables plus an intercept term: the percentile rank of ICU occupancy, number of current ED patients, and a dummy for whether the patient is in low or high severity class. This simplified model uses the same information as in our dynamic discrete choice model. Table OS.9 reports the estimated coefficient γ_{ICU} of ICU occupancy for the ICU admission decision ($d_{it} = \text{ICUAdm}$). The results for the full (resp. simplified) multinomial logistic model are given in the second (resp. third) columns respectively. We see that the results are qualitatively similar to that for all hospitals combined in Table EC.1. Specifically, higher ICU occupancy is still associated with decreased likelihood of ICU admission decisions for most hospitals, although some coefficients are not statistically significant due to a much smaller sample size. Moreover, the identified discount factors $\hat{\beta}$ and the coefficients of ICU occupancy γ_{ICU} are highly negatively correlated: the correlation is -0.768 and -0.774 for the full and simplified model respectively. This means that the dynamic discrete choice model and the multinomial logistic model provide consistent evidence for hospitals' discounting behaviors. In the full multinomial logistic model (EC.4), the McFaddean's R^2 for individual hospital regressions ranges from 0.14 to 0.24, with an average of 0.17. This level is also similar to our dynamic discrete choice model.

EC.4. Individual Hospital Results

Table EC.2 provides the full estimate results by individual hospital for our main model.

EC.5. Discussion on Model Identification

Action utility and customer classes: As described in Section 2, the parameters $\{u_{w,\iota}, u_{a,\iota}, u_{r,\iota}\}$ represent the average action utility within each customer class. Then, the term $\varepsilon(d_t)$ in (3) captures the idiosyncratic utility shock related to specific customer or action. Such set-up is common in dynamic discrete choice models (see, e.g., Gordon 2009, Komarova et al. 2018).

Due to the curse of dimensionality, it is prohibitive in our model to let the action utility depend on each customer's characteristics (e.g., patient's severity). The curse of dimensionality is a common challenge for the estimation of dynamic discrete choice models, as one usually needs to solve the dynamic programming equations for every candidate parameter (Rust 1987). This issue is notably significant in our context due to the fact that all the customers use the same FSU with limited capacity. Thus, our state space must include the information of *all* customers in the GK. With L customer classes, the dimension of state space is

$$\text{Dim of state space} = \prod_{i=1}^{L}(Q_i + 1) \times (B + 1),$$

where Q_i is the GK capacity for ith customer class and B is the FSU capacity. Clearly, the dimension of state space increases exponentially fast with the number of customer classes. For these reasons, in our main setting, we include two customer classes to capture the heterogeneity in their decisions and use utility parameters to represent the average utilities for each class. In our empirical study on ICU admission, the dimension of state space in our model is already large: the average dimension is 1318 for the 22 hospitals, with a maximum of 3584. This presents a significant challenge to the model estimation, especially since we are identifying the discount factor jointly with the utility parameters.

From the modeling perspective, our set-up is similar to the consumer purchasing/replacement problems studied in Gordon (2009) and Gowrisankaran and Rysman (2012), in which the dynamic decision is made at the system-level. Accordingly, the model state has to include the characteristics (e.g., price, quality) of *all* available products in the market. In both papers, important simplifications and assumptions have to be made

Table EC.2 Estimation results of structural model by individual hospital

Hosp	Num. of periods	Discount Factor $\hat{\beta}$	Low Severity $\hat{u}_{w,l}$	Low Severity $\hat{u}_{r,l}$	High Severity $\hat{u}_{w,h}$	High Severity $\hat{u}_{r,h}$	R^2
1	8016	0.3	−0.015	1.490	−0.749	0.301	0.19
		(0.025)	(0.016)	(0.031)	(0.060)	(0.041)	
2	6012	0.5	−0.001	1.465	−0.659	0.519	0.23
		(0.026)	(0.038)	(0.028)	(0.054)	(0.055)	
3	8016	0.1	−0.124	2.558	−1.560	1.185	0.08
		(0.052)	(0.157)	(0.052)	(0.155)	(0.067)	
4	8016	0.4	−0.065	2.307	−0.940	0.798	0.15
		(0.011)	(0.034)	(0.040)	(0.053)	(0.044)	
5	6924	0.1	−0.179	1.786	−1.147	0.336	0.10
		(0.036)	(0.096)	(0.053)	(0.117)	(0.076)	
6	8016	0.2	−0.077	1.904	−0.970	0.531	0.14
		(0.041)	(0.039)	(0.045)	(0.080)	(0.048)	
7	6948	0.9	−1.755	2.687	−1.876	1.453	0.10
		(0.177)	(0.515)	(0.083)	(0.328)	(0.096)	
8	7848	0.1	−0.249	1.446	−0.764	0.380	0.15
		(0.062)	(0.111)	(0.034)	(0.089)	(0.051)	
9	6180	0.5	0.000	1.745	−0.789	0.707	0.18
		(0.035)	(0.036)	(0.040)	(0.062)	(0.050)	
10	7320	0.1	−0.629	2.339	−1.630	0.929	0.07
		(0.325)	(0.800)	(0.053)	(0.466)	(0.086)	
11	8016	0.9	−1.127	1.774	−1.366	0.104	0.08
		(0.248)	(0.538)	(0.064)	(0.274)	(0.108)	
12	4668	0.6	−0.099	2.121	−0.558	0.637	0.22
		(0.014)	(0.033)	(0.051)	(0.058)	(0.071)	
13	8016	0.1	−0.044	2.229	−1.260	0.785	0.13
		(0.016)	(0.051)	(0.045)	(0.110)	(0.055)	
14	8016	0.3	−0.147	1.951	−0.956	0.662	0.19
		(0.050)	(0.063)	(0.036)	(0.082)	(0.044)	
15	6912	0.1	−1.082	2.155	−2.275	0.762	0.10
		(0.071)	(0.180)	(0.045)	(0.219)	(0.076)	
16	8016	0.1	−0.199	2.319	−1.233	1.005	0.11
		(0.000)	(0.072)	(0.050)	(0.123)	(0.066)	
17	6588	0.9	−1.669	2.065	−1.530	0.660	0.09
		(0.082)	(0.202)	(0.069)	(0.134)	(0.081)	
18	8016	0.4	−0.154	2.358	−1.391	1.005	0.04
		(0.046)	(0.065)	(0.060)	(0.114)	(0.069)	
19	6576	0.1	−0.886	1.809	−1.885	0.517	0.12
		(0.000)	(0.069)	(0.039)	(0.138)	(0.056)	
20	8004	0.3	−0.116	2.260	−1.335	0.862	0.09
		(0.043)	(0.054)	(0.063)	(0.131)	(0.084)	
21	4008	0.7	−0.135	2.002	−0.612	0.240	0.12
		(0.047)	(0.058)	(0.094)	(0.096)	(0.120)	
22	4008	0.3	0.000	1.513	−0.979	0.175	0.09
		(0.042)	(0.030)	(0.068)	(0.130)	(0.111)	

Note: The second column reports the number of two-hour periods in each hospital, and the last column provides the McFadden's pseudo R^2. Standard errors are reported in parenthesis. They are estimated by bootstrapping with 500 trials. We find that the standard errors of the estimated $\hat{\beta}$'s are smaller than 0.05 (resp. 0.1) for 15 (resp. 19) out of the 22 hospitals. Although not reported here, in the bootstrapping, we find that the estimated $\hat{\beta}$ remains the same or only changes to the adjacent levels (plus or minus 0.1) in more than 90% of the trials for 20 of the 22 hospitals (except Hospitals 14 and 15).

to mitigate the curse of dimension of state space. For example, Gordon (2009) only considers two product types for each of the two firms in their model, by aggregating the product characteristics within each type. This is similar to our setting with two customer classes. In contrast, our set-up is different from the structural models in, e.g., Emadi and Staats (2020) and Yu et al. (2016), in which the dynamic decision is modeled at the *individual* level. Thus, the state space only needs to contain the information for each individual. This substantially reduces the computational burden and facilitates the use of parametric or random utility coefficients to account for agent heterogeneity. Moreover, we identify the discount factor and action utilities jointly in the dynamic discrete choice model, which brings additional technical challenge.

Counterexample for non-identification: The key intuition for model identification is that there is enough variation in the observed states and actions relative to the parameter space. We now provide a simple counter-example to show that the model cannot be identified when there is not sufficient variation in the states and actions.

We consider an FSU with only one bed. We assume there is always one patient in the GK in every period. The system can take three actions for the patient: keep her waiting in GK, admit her to SSU (with unlimited capacity), or admit her to FSU. The last action is only feasible when FSU is empty. After one action is taken in each period, the FSU state may transit due to external arrival or patient departure. The transition probabilities are given by $P = [p_{00}, p_{01}; p_{10}, p_{11}]$, where p_{ij} is the probability that the FSU state changes from i to j. The action utilities are given by u_w for GK waiting, u_r for SSU admission, and u_a for FSU admission. We assume an additive idiosyncratic term in the utility as in (3). The discount factor is $\beta \in (0,1)$. Denote the action probabilities by $f(\iota_1 | n_t^I = 0)$ for $\iota_1 \in \{w, r, a\}$ and $f(\iota_2 | n_t^I = 1)$ for $\iota_2 \in \{w, r\}$. These action probabilities can be computed similarly to eq. (9) in our main setting.

We perform the following numerical study. First, we compute the correct choice probabilities $f(\iota_1 | n_t^I = 0)$ and $f(\iota_2 | n_t^I = 1)$ under a given set of parameters (β, u_a, u_r, u_w). Then, we assume an alternative, incorrect discount factor β' and search for the corresponding utility parameters (u_a', u_r', u_w') to fit the correct choice probabilities. The FSU admission utility is normalized to zero ($u_a = 0$).

We find that the discount factor cannot be identified in this model. In particular, for an incorrect discount factor β', we can always find the utility parameters (u_w', u_r') such that the choice probabilities coincide with the true values for every state and action. Thus, the model is unidentified due to observational equivalence. In particular, the choice probabilities can be equivalently rationalized by a small discount factor combined with high action utilities or a large discount factor combined with low action utilities (see Figure OS.3 of the Online Supplement). This counter-example shows that we need enough variation in the observed states and actions to identify the discount factor and utility parameters jointly.

State-dependent action set: A novelty of our identification is that our model has state-dependent action set, which differs from Komarova et al. (2018). We now conduct a simulation study based on the ICU admission problem to show that using state-dependent action sets is essential to correctly identify the model.

We first simulate a series of state and action pairs using the correct model with state-dependent action sets for a given set of discount factor and utility parameters. We define the action in period t as $d_t = (a_{l,t}, r_{l,t}, a_{h,t}, r_{h,t})$, which includes the numbers of ICU and non-ICU admissions for the low and high severity classes, respectively. We then select a state $\bar{s} = (n_l^G, n_h^G, n^I)$ as $n_\iota^G = \max\{n_{\iota,t}^G\}$ for $\iota \in \{l, h\}$ and $n^I = B - n_l^G - n_h^G$. That is, $\bar{s}$ has the maximum numbers of ED patients for both severity classes in the simulation and enough empty beds in the ICU for admitting all ED patients. It is easy to verify that with the state $\bar{s}$, all simulated actions $d_t = (a_{l,t}, r_{l,t}, a_{h,t}, r_{h,t})$ are feasible.

We then estimate the model parameters using the simulated data from the misspecified model without state-dependent action sets. For a candidate set of parameters $\theta = (\beta, u_{w,l}, u_{r,l}, u_{w,h}, u_{r,h})$, we solve the value function using the correct model with state-dependent action sets. However, in the estimation, we incorrectly assume that the system state is given by $\bar{s}$ in every period t. That is, we use the choice probability $f(d_t|\bar{s})$ and compute the likelihood function as $l^f = \Pi_{t=1}^T f(d_t|\bar{s})$ for the simulated data. We search for the discount factor and utility parameters which maximize the likelihood function. We use the same normalization as in main specification and restrict the discount factor to the grid $\{0.1, 0.2, ..., 0.9\}$.

We select three representative hospitals, which have small, medium, and large estimated discount factors under our main specification, to conduct the simulation study. For each hospital, we generate 100 paths of state-action pairs using the estimated parameters (see Table EC.2). Each path includes 20,000 two-hour intervals with a one-month warm-up period. For each path, we estimate the discount factor and utility parameters using the correct state-dependent action set $\Pi(s_t)$ as in our main specification, or the incorrect state-independent action set $\Pi(\bar{s})$ as described above. For each estimation approach, we take the average over the 100 paths to obtain the final estimates. The results are reported in Table OS.10.

We find that when the state-independent action set $\Pi(\bar{s})$ is used in identification and estimation, the estimated discount factor is always 0.9 and the utility parameters are largely biased from their true values (see the "State-independent action set $\Pi(\bar{s})$" columns). In particular, the estimated waiting utilities $\hat{u}_{w,l}$ and $\hat{u}_{w,h}$ are close to zero with an upward bias in all cases, and the non-ICU admission utilities $\hat{u}_{r,l}$ and $\hat{u}_{r,h}$ are largely underestimated with a substantial downward bias. In contrast, using the correct state-dependent action set $\Pi(s_t)$ produces estimated parameters (both discount factor and and utility parameters) that are close to the true values (see the "Correct action set" columns).

The estimation bias can be interpreted as follows. In the sample, the hospitals slow down ICU admissions when the ICU becomes full, as hospitals would attempt save ICU capacity for future patients. However, under the state $\bar{s}$, it is always feasible to admit all ED patients to the ICU. Thus, the hospitals' capacity saving behaviors of keeping patients in the ED can only be (partially) rationalized by a large discount factor, a high (non-negative) waiting utility, and/or a low non-ICU admission utility. This leads to the observed bias in the estimation results. The above comparison shows that it is essential to account for the state-dependent action set in the identification and estimation of dynamic discrete choice models.

EC.6. Capacity Intervention of Adding One ICU Bed

We first evaluate the counterfactual impact of adding one more bed in ICU, while keeping the arrival and service rates fixed. For the hospitals in our study, adding one ICU bed has limited impact on the average ICU occupancy level. We focus on the impact on high ICU congestion measures introduced in (14) – (17).

Table EC.3 reports the reductions in ICU congestion measures for six selected hospitals. We see that adding one bed in the ICU indeed leads to substantial reductions in ICU congestion. For some hospitals, such reduction translates to 18 fewer days and 50 fewer patients exposed to high ICU congestion ("ΔDays HighCgstn" and "ΔPats HighCgstn" columns), as well as 10 more external arrivals admitted ("ΔPats Balk" column) in each year. The relative drops in $\Pr(\text{HighCgstn})$ and $\Pr(\text{Balk})$ are greater than 30% in most cases. As patients who require ICU care are usually very unstable, the observed effects have the potential to cause great improvements in their medical outcomes. Although not reported here, the standard deviations computed from the simulation confirm that all differences are statistically significant at the 1% level.

Table EC.3 Counterfactual Estimates of Impact when Adding One Bed in ICU (Select Hospitals)

Hosp	$\hat{\beta}$	B	ΔDays HighCgstn (in # days)	ΔPats HighCgstn (in # patients)	Relative drop $\Pr(\text{HighCgstn})$	ΔPats Balk (in # patients)	Relative drop $\Pr(\text{Balk})$
1	0.3	21	6.27	52.26	35.3%	6.73	38.8%
2	0.5	26	6.78	52.78	24.9%	8.43	27.5%
5	0.1	11	18.74	68.04	38.2%	8.84	39.0%
8	0.1	16	10.90	64.88	37.5%	7.92	41.5%
9	0.5	22	8.60	78.61	31.7%	13.57	35.3%
14	0.3	36	2.36	25.90	24.0%	5.26	27.6%

Next, we look into how the benefit of adding one bed in ICU varies across hospitals. Table EC.4 reports the correlations of the reductions in the three congestion measures ("ΔDays HighCgstn", "ΔPats HighCgstn", and "ΔPats Balk" columns in Table EC.3) with the average ICU occupancy levels and the estimated discount factors across the 22 hospitals in our sample. By the first two columns, all three effects are positively correlated with both ICU occupancy levels and estimated discount factors. This relationship still holds when we perform the following linear regression for the three effects using the estimates of the 22 hospitals:

$$\Delta\text{ICUCgstn}_i = \lambda_0 + \lambda_1 \hat{\beta}_i + \lambda_2 \text{ICUOccu}_i + \varepsilon_i.$$

By the last two columns of Table EC.4, the benefit of adding an ICU bed is positively associated with both ICU occupancy and the estimated discount factor, though the relative importance of the two varies across ICU congestion metrics. These results suggest that the benefit of adding one bed in the ICU tends to be greater for ICUs with higher occupancy levels and/or hospitals that account more for the system-level impact of their admission decisions. Although we do not interpret these results as prescriptive, they illustrate that the balancing behavior of hospitals is an important factor in capacity planning decisions. Notably, the discount factor reflects the how the hospital balances the current patients and future system capacity, while the ICU occupancy measures the average system utilization. Thus, they capture different aspects in the ICU admission decisions and are complementary in capacity planning. This reveals the benefit of understanding the hospitals' discounting behaviors from the data.

Table EC.4 Correlations between Impact of Adding One ICU Bed and ICU Occupancy or $\hat{\beta}$

	Correlation		Regression coefficient	
	$\hat{\beta}$	ICUOccu	$\hat{\beta}$	ICUOccu
ΔDays HighCgstn	0.69^{***}	0.36	18.58^{**}	3.98
ΔPats HighCgstn	0.53^{*}	0.77^{**}	19.64	140.90^{***}
ΔPats Balk	0.60^{**}	0.75^{**}	5.72^{*}	26.55^{**}

Note: $^{*}p < 0.05$, $^{**}p < 0.01$, and $^{***}p < 0.001$.

EC.7. Robustness Checks

In this section, we conduct several robustness checks for our empirical results in Section 4. Our main findings are largely robust in these settings.

Choice of Normalization Levels: We perform two robustness checks for the normalization choice of $(u_{a,l}, u_{a,h})$, which can affect the identification results (see Section 2.2). First, we estimate our model using the periods with same or similar ED states, which mitigates the impact of normalization choice by Proposition 4. We select the state to be $(n_l^G, n_h^G) = (2, 1)$, i.e., with two low severity patients and one high severity patient. This state includes multiple ED patients with both severity classes. Thus, the action space is large enough to reflect the balancing behavior of hospitals. In addition, the occurrence of the state is not too rare, so the model can be effectively estimated by combining all hospitals. In the estimation, we now calculate the log-likelihood l^f in (13) only using the actions under the ED state $(n_l^G, n_h^G) = (2, 1)$.

The results are reported in Table EC.5 for all hospitals combined. The first row repeats Table 1 for our original model with all ED states. The second row reports the estimation results with only the ED state $(n_l^G, n_h^G) = (2, 1)$. As an additional check, we report in the third row the estimated parameters when we use the ED states $(n_l^G, n_h^G) = (2, 1)$ or $(1, 1)$. This further increases the sample size. The ICU admission utilities are assumed to be zero in the estimations. We see that $\hat{\beta}$ equals to 0.3 in all the three cases, suggesting the identified discount factor is not affected when we limit to the select ED states. Figure OS.4 further shows that the likelihood decreases monotonically as we deviate from the optimal $\hat{\beta} = 0.3$ in all three cases. In addition, the waiting and non-ICU admission utilities are also similar to the original model for both classes. These findings support the robustness of our identification results.

As discussed in Section 3.3, we expect the utilities of ICU admission to be close to zero in our empirical setting. In the second robustness check, we further estimate our model for each hospital separately under

Table EC.5 Estimation Results with Different ED States (All Hospitals Combined)

(n_l^E, n_h^E)	Size	$\hat{\beta}$	$\hat{u}_{w,l}$	$\hat{u}_{r,l}$	$\hat{u}_{w,h}$	$\hat{u}_{r,h}$
All	154,140	0.3 (0.005)	-0.071 (0.001)	1.950 (0.010)	-0.932 (0.018)	0.671 (0.013)
(2,1)	4,230	0.3 (0.019)	-0.013 (0.005)	2.001 (0.042)	-0.966 (0.051)	0.775 (0.036)
(2,1) or (1,1)	11,706	0.3 (0.023)	-0.071 (0.044)	2.215 (0.031)	-1.059 (0.040)	0.727 (0.022)

non-zero ICU admission utilities $u_{a,h}$ and $u_{a,l}$. Here all ED states are used in the estimation. In particular, we estimate the model with different normalization levels for the ICU admission utilities. We consider the cases with both $u_{a,h} = u_{a,l}$ and $u_{a,l} < u_{a,h}$ (we assume the ICU admission utility is higher for the high-severity class). The values of $u_{a,h}$ and $u_{a,l}$ are varied from -0.3 to 0.3. As in our main setting, we still assume that the ED admission utility is lower than the utility of ICU and non-ICU admission.

The estimated discount factors for individual hospitals are shown in Table OS.11. Each column in the table corresponds to a normalization choice $(u_{a,l}, u_{a,h})$. We find that the estimated discount factors are very robust under alternative normalization choices. The estimated discount factor $\hat{\beta}$ remains unchanged or only changes to adjacent levels (increase or decrease by 0.1) in all cases. We also compute the correlation (across hospitals) between $\hat{\beta}$ in the original model and each of the alternative settings. The correlation is higher than 0.9 in all cases. Thus, our results are robust to the choice of normalization levels.

Stratified Sample and Time Trend: As a robustness check, we also estimate the structural model with stratified samples based on flu vs non-flu seasons, as well as day and night periods. The flu season ranges from November to March, and the non-flu season includes the rest. The day periods include the twelve hours from 7AM to 7PM. We re-estimate the ED arrival rates $\lambda_{Q,l}$ and $\lambda_{Q,h}$, external ICU arrival rate λ_E, and ICU departure rate μ_I for each stratified sample. The ED and ICU capacities Q_l, Q_h, and B are set as those for the full sample in Table OS.6. We estimate the structural model with the stratified samples for all hospitals combined and for individual hospitals separately.

Table EC.6 reports the results from the combined estimation. We can see the estimated parameters are close to that in our main specification (first two rows). The estimated discount factor is 0.4 for the flu season, 0.2 for the non-flu season, and 0.4 for the day sample, and 0.1 for the night sample. The lower discounting level during the night can be explained as there may be more medical and operational constraints that hinder the hospital from responding to the system state during the night (e.g., shortage in nurses). For individual hospitals, we find the heterogeneity in discount factors estimated from the stratified data is largely consistent with that from the full sample given in Table EC.2. In particular, the across-hospital correlations between the estimated $\hat{\beta}$ from stratified and full sample are significantly positive: 0.87 for the flu season, 0.67 for the non-flu season and day sample, and 0.46 for the night sample. Thus, our main findings from the structural model are robust to accounting for these potential temporal variations.

In Section OS.4 of the Online Supplement, we further develop and estimate a structural model that explicitly accounts for time trend in the system transition during a day. The model includes a time indicator for each two-hour interval in a day in as a state variable, and allow the arrival and departure rates to vary with time. Our main findings remain very robust after the time trend is included in the model.

Table EC.6 Estimation of Structural Model with Stratified Data: All Hospitals Combined

	Size N	Discount factor $\hat{\beta}$	Low Severity $\hat{u}_{w,l}$	Low Severity $\hat{u}_{r,l}$	High Severity $\hat{u}_{w,h}$	High Severity $\hat{u}_{r,h}$	R^2
All	154,140	0.3	−0.071	1.950	−0.932	0.671	0.14
		(0.005)	(0.001)	(0.010)	(0.018)	(0.013)	
Flu	60,031	0.4	−0.028	1.936	−0.755	0.674	0.14
		(0.004)	(0.010)	(0.012)	(0.021)	(0.019)	
Non-flu	108,737	0.2	−0.058	1.960	−1.086	0.670	0.13
		(0.005)	(0.010)	(0.009)	(0.023)	(0.016)	
Day	84,372	0.4	−0.001	2.006	−0.806	0.655	0.13
		(0.004)	(0.009)	(0.011)	(0.020)	(0.018)	
Night	84,396	0.1	−0.045	1.875	−0.971	0.684	0.14
		(0.005)	(0.011)	(0.009)	(0.024)	(0.016)	

Standard error is reported in parenthesis.

Finer Grid Search, Three Patient Class, and Change in Unit Capacities: For computational purposes, we conduct a grid search over the discount factor with a step of 0.1 in our main specification. This facilitates the numerical implementation of our estimation as the state space in our model is relatively large. As a robustness check, we have also done a finer grid search with 0.01 increment for select Hospitals 4 and 9. The identified β's from the finer grid are very close to those from the original grid (see Figure OS.5 of the Online Supplement).

In Section OS.5 of the Online Supplement, we develop and estimate a dynamic structure model with three patient classes for all hospitals combined and each individual hospital. The estimation results (discount factor and utility parameters) are very similar under the two-class and three-class models (see Tables OS.4 and OS.13). Finally, we estimate our model under alternative ICU and ED capacity levels, which are set according to the data (see Section EC.2.2). In robustness checks, we increase or decrease the ICU capacity B by one bed, as well as drop the buffer term $\lfloor M_{A,i} \rfloor$ in the ED capacities Q_i for $i \in \{l, h\}$. The identified discount factors are reported in Tables OS.14 and OS.15, which are very similar to our main specification. These results further demonstrate the robustness of our findings.

EC.8. Summary of Model Assumptions

In this section, we present the major assumptions used in our empirical study as well as their economic justifications and robustness checks. These discussions are summarized in Table OS.16.

Linear utility function: As discussed in the paper, we achieve joint identification of discount factor and utility parameters by leveraging the results in Komarova et al. (2018), which require per-period payoff to be linear in utility parameters. This assumption is critical for our identification. It cannot be tested from data, but we believe it is a reasonable assumption to capture the first-order effect in our model. In particular, the hospitals' per-period (perceived) utility should be the sum of action utility for each patient affected by the decision, which is reflected by the linear utility function in (2). In addition, as shown in Table OS.18, the number of ED patients in each period is relatively small, mitigating the non-linearity concern regarding. Besides, we also find that the ED waiting time is relatively short for most patients (see Table OS.17), suggesting that patient deterioration in ED to have little impact.

Constant utility parameters:We assume the utility parameters are constant and independent of the state, which is standard in dynamic discrete choice model literature. This assumption is crucial for our identification and interpretation of the discount factor in our model. As discussed in Section 4.2, we expect the benefit from ICU care to remain relatively stable even when ICU is congested, since ICU typically allocates dedicated resource to each bed and maintains a fixed nurse-to-bed ratio. In addition, we find no significant impact of ICU occupancy on patient outcomes (see Table OS.2 in Section OS.3). These results provide suggestive evidence that our main findings are not driven by congestion-dependent ICU admission utility.

Normalization level: In our main model, we normalize the ICU admission utilities of both classes to be zero. The normalization assumption is necessary to achieve identification, and the specific normalization choice can affect our estimates (e.g., Aguirregabiria and Mira 2010). The normalization choice cannot be tested from the data, but we justify it based on our empirical context and provide multiple robustness checks.

The normalization assumption is needed for identification in most dynamic discrete choice models, as we cannot identify the utilities of all actions from data (e.g., Rust 1994, Magnac and Thesmar 2002). In addition, setting zero utility for a reference action is the most common practice (e.g., Magnac and Thesmar 2002, Ryan 2012, Aguirregabiria and Mira 2010). In general dynamic discrete choice models, the normalization choice can affect the identification results and counterfactual estimates (see e.g., Magnac and Thesmar 2002, Norets and Tang 2014, Aguirregabiria and Suzuki 2014, and Abbring and Daljord 2020).

As discussed in Section 3.3, the assumption of zero ICU admission utility is based on our empirical context. First, setting ICU admission utility to zero provides a proper reference point for perceived utility across all ED states, as ICU admission is the best action hospitals can do in terms of patient's clinical outcomes. Second, the marginal utility of admitting a patient to ICU should be small due to its very high fixed and operating cost. Due to data limitation, we assume the ICU admission utility is the same for the two severity classes. Note that if a patient is eventually admitted to ICU, it is likely because his unobserved factors or latest conditions imply significant risk, regardless of the pre-assigned severity class. From the data, we find that although the high severity class is four times more likely to be admitted to ICU than the low severity class (34% vs. 8%), the difference in their ICU length-of-stay (LOS) is much smaller. As shown in Table OS.19, the average ICU LOS of the low severity class is roughly 70% of that of the high severity class. In addition, the ICU LOS exhibits substantial variation for both severity classes.

Given the normalization choice impacts our identification results, we have performed multiple robustness checks using stratified samples (periods with same ED state) and alternative normalization levels. They are discussed in Section EC.7. We show that our main findings remain robust in these tests. Moreover, under current normalization assumption, the estimated dynamic discrete choice model fits our observed data well (Section OS.2) and provide evidence for hospitals' discounting behaviors that is consistent with the multinomial logistic model (Section EC.3).

Lower ED waiting utility: We assume the utility of keeping patients waiting in ED is lower than that of admitting them to the ICU and non-ICU units. This assumption is important for our empirical study, as it rules out unrealistic model estimates that are improper in our setting.

The assumption cannot be directly tested from data, but it is justified based on our empirical context. In our setting, the ED waiting time is defined as the time difference between the time the order to admit is made and the time the patient is eventually admitted to an inpatient unit. Thus, ED waiting is generally detrimental to our patient cohort, given they must be admitted to the hospital. It represents an additional cost in the patient rerouting process and is only a temporary state due to frictions or constraints.

From the data, we indeed find that the ED waiting time is relatively short (see Table OS.17) for most patients, suggesting that ED waiting is on average an undesirable state in patient flow. Therefore, we assume the ED waiting should bring lower perceived utility than ICU and non-ICU admission. Note that the utility parameters represent the *average* perceived utility associated with each action for the two patient classes. The patient-specific deviations from the average utilities are captured by the idiosyncratic term, as standard in dynamic discrete choice models (Rust 1987, Komarova et al. 2018).

Patient classes: We divide patients to two severity classes in our model and use six deterministic parameters to represent the average perceived utilities for their actions. As discussed in Section EC.5, this is necessary for our identification as our dimension of state space grows fast with the number of patient classes. We believe this assumption is innocuous as two severity classes are sufficient to capture the heterogeneity in the patients' need for service. As a robustness check on this assumption, we have shown that our main findings remain very similar when we define three patient classes with low, medium, and high severity levels. This is described in Section OS.5 (Tables OS.4 and OS.13).

Other model primitives: We assume constant model primitives including ICU and ED capacities, arrival rates to ED and ICU, and ICU departure rate. These set-ups are in line with the dynamic discrete choice model literature (e.g., Rust 1994). The model parameters are directly estimated from the observed data using standard methods. We conduct multiple robustness checks on these model primitives. For ICU and ED capacities, we estimate the model after increasing or decreasing the ICU capacity by one bed from the current level, as well as after dropping the buffer term in the ED capacity. The estimation results remain very similar to our main specification for both all hospitals combined and individual hospitals (see Tables OS.14 and OS.15). For arrival and departure rates, we estimate the model after incorporating the time trend in the state variable (see Section OS.4). That is, we allow the arrival and departure rates to vary for each two-hour interval. In addition, we estimate the model separately for daytime and nighttime, as well as for flu and non-flu seasons. Our main findings remain robust in the above settings.

In line with the literature, we assume the arrivals per period follow Poisson distributions. To validate this, we compute the ratio between the variance of arrivals per period to the mean of arrivals for each hospital. As reported in Table OS.20, the ratio is close to one for most hospitals, consistent with the Poisson distribution. We also assume the number of ICU departures follows a binomial distribution in each period with a constant rate. Table OS.21 reports the estimated departure rates for ICU occupancy levels in each decile bin (all hospitals combined). We see that the departure rates are relatively stable for different decile bins.

EC.9. Impact of Misspecifying Discount Factor in Identification

In this section, we use counterfactual simulations to show how pre-specifying an "incorrect" discount factor leads to deviations in evaluating the high ICU congestion periods. We consider four choices of discount factors for each hospital: $\beta = 0$ for a pure myopic model, $\beta = 0.3$ as the average level from the combined estimation, $\beta = \hat{\beta}$ at the estimated level, and $\beta = 0.99$ as is commonly used in the literature. For a given discount factor, the utility parameters are estimated accordingly for each hospital. By our estimation procedure, the estimated discount factor leads to the highest likelihood of hospitals' observed actions (see Figures OS.4 and OS.5), i.e., it fits the data better than alternative discount factors. Thus, we use it as the benchmark for comparison. We also perform the likelihood-ratio test to compare different models. For all hospitals combined, the models with $\beta = 0$ and $\beta = 0.99$ are rejected compared to the identified model at the 0.1% significance level. For individual hospitals, the model with $\beta = 0$ (resp. with $\beta = 0.99$) is rejected compared to the identified model at the 1% level for 16 (resp. 15) out of the 22 hospitals.

We use counterfactual simulations to estimate the ICU congestion under three scenarios. The first is the base case with the estimated system set-up. We then consider two system changes. In the first one, we increase the ED arrival rates $\lambda_{Q,l}$ and $\lambda_{Q,h}$ by 50%. In the second one, we keep the ED arrivals at the original level but reduce the ICU capacity B by one bed. The two system changes reflect scenarios when ICU capacity gets constrained (e.g., during a pandemic). Correctly evaluating the impact of such shocks is crucial to hospitals for managing their ICUs.

Table EC.7 provides the estimates of total numbers of patients exposed to high ICU congestion and external arrivals who balk under different choices of discount factor (with corresponding utility estimates) and the three scenarios of system set-up. The numbers in the table represent the sum of the 22 hospitals in our sample. The top row "$\beta = \hat{\beta}$" shows the estimates when we jointly identify discount factor and utility parameters for each hospital. Thus, we use it as the benchmark for comparison. For other choices of β, their relative deviations compared with $\beta = \hat{\beta}$ are reported in the parenthesis.

Table EC.7 Estimated ICU Congestion Metrics under Difference Choices of Discount Factors and Corresponding Utility Estimates

Discount factor	Base case		Increasing ED arrival by 50%		Reducing one ICU bed	
	Pats HighCgstn	Pats Balk	Pats HighCgstn	Pats Balk	Pats HighCgstn	Pats Balk
$\beta=\hat{\beta}$	2,356	352	4,408	678	3308	546
$\beta=0$	4,319 (83.4%)	718 (103.8%)	7,384 (67.5%)	1,282 (89.0%)	5,445 (64.6%)	995 (82.5%)
$\beta=0.3$	2,961 (25.7%)	461 (30.9%)	5,335 (21.0%)	866 (27.6%)	3,993 (20.7%)	692 (26.8%)
$\beta=0.99$	2,223 (−5.6%)	324 (−8.1%)	3,718 (−15.6%)	554 (−18.3%)	3,065 (−7.3%)	500 (−8.3%)

We have the following observations. First, setting $\beta = 0$ or 0.3 in the structural model substantially overestimates the ICU congestion compared with using the identified $\hat{\beta}$. In the base case, setting $\beta = 0$ (resp. $\beta = 0.3$) overestimates the number of patients under high ICU congestion by 83.4% (resp. 25.7%) and the external arrivals who balk by 103.8% (resp. 30.9%). In contrast, assuming $\beta = 0.99$ would underestimate the ICU congestion. In the base case, it underestimates the patients under high ICU congestion (resp. externals who balk) by 5.6% (resp. 8.1%). Such deviation becomes substantially larger when the ICU gets more constrained in adverse scenarios. After we increase ED arrivals by 50% (resp. reduce one ICU bed), the relative underestimation with $\beta = 0.99$ increases to 15.6% (resp. 7.3%) for high ICU congestion patients and 18.3% (resp. 8.3%) for balking external arrivals. The above findings are in line with our interpretation of the discount factor: Hospitals with larger discount factors respond to ICU state more proactively when making admission decisions, i.e., reducing ICU admission likelihood when ICU occupancy is high. This mitigates the high ICU congestion periods.

We also compute other key system statistics under different discount factors, including overall ICU admission probabilities, ICU occupancy, and ED waiting times. We find that using $\beta = 0.99$ and $\beta = \hat{\beta}$ with their respective estimated utility parameters in the model produce close estimates for these system statistics. This reflects the potential "observational equivalence" issue that hinders the identification (e.g., Komarova et al. 2018). However, as shown in Table EC.7, the estimates of high ICU congestion clearly hinge on the choice of discount factor. This is because the impact of hospital's balancing behavior is more visible during high ICU congestion periods. The above finding highlights the importance of identifying the discount factor using observed data. Similar analysis is conducted in Ching and Osborne (2020) under a different setting. It shows that using an "incorrect" discount factor of 0.995 would largely overestimate the increase in sales from price promotions.

References

Abbring, Jaap H, Øystein Daljord. 2020. Identifying the discount factor in dynamic discrete choice models. *Quantitative Economics* **11**(2) 471–501.

Aguirregabiria, Victor, Pedro Mira. 2010. Dynamic discrete choice structural models: A survey. *Journal of Econometrics* **156**(1) 38–67.

Aguirregabiria, Victor, Junichi Suzuki. 2014. Identification and counterfactuals in dynamic models of market entry and exit. *Quantitative Marketing and Economics* **12**(3) 267–304.

Ching, Andrew T, Matthew Osborne. 2020. Identification and estimation of forward-looking behavior: The case of consumer stockpiling. *Marketing Science* .

Emadi, Seyed Morteza, Bradley R Staats. 2020. A structural estimation approach to study agent attrition. *Management Science* **66**(9) 4071–4095.

Gordon, Brett R. 2009. A dynamic model of consumer replacement cycles in the pc processor industry. *Marketing Science* **28**(5) 846–867.

Gowrisankaran, Gautam, Marc Rysman. 2012. Dynamics of consumer demand for new durable goods. *Journal of Political Economy* **120**(6) 1173–1219.

Komarova, Tatiana, Fabio Sanches, Daniel Silva Junior, Sorawoot Srisuma. 2018. Joint analysis of the discount factor and payoff parameters in dynamic discrete choice models. *Quantitative Economics* **9**(3) 1153–1194.

Magnac, Thierry, David Thesmar. 2002. Identifying dynamic discrete decision processes. *Econometrica* **70**(2) 801–816.

Norets, Andriy, Xun Tang. 2014. Semiparametric inference in dynamic binary choice models. *Review of Economic Studies* **81**(3) 1229–1262.

Rust, John. 1987. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society* 999–1033.

Rust, John. 1994. Structural estimation of markov decision processes. *Handbook of Econometrics* **4** 3081–3143.

Ryan, Stephen P. 2012. The costs of environmental regulation in a concentrated industry. *Econometrica* **80**(3) 1019–1061.

Yu, Qiuping, Gad Allon, Achal Bassamboo. 2016. How do delay announcements shape customer behavior? an empirical study. *Management Science* **63**(1) 1–20.

Online Supplement to "Structural Estimation of Intertemporal Externalities with Application in ICU Admissions"

Yiwen Shen
Department of ISOM, HKUST, yiwenshen@ust.hk

Carri W. Chan
Decision, Risk, and Operations, Columbia Business School, cwchan@columbia.edu

Fanyin Zheng
Decision, Risk, and Operations, Columbia Business School, fanyin.zheng@columbia.edu

Gabriel Escobar
Division of Research, Kaiser Permanente, gabriel.escobar@kp.org

OS.1. Proofs

OS.1.1. Proof for Proposition 2

By Rust (1987), the value function $\tilde{V}(s)$ can be expressed as:

$$\begin{aligned}\tilde{V}(s) &= \sum_{s'} \int_{\varepsilon'} V(s', \varepsilon')\, g(s'|s) q(\varepsilon'|s') d\varepsilon' \\ &= \sum_{s_t} \int_{\varepsilon_t} \sup_{d_t \in \Pi(s_t)} \mathsf{E} \left\{ \sum_{j=t}^{\infty} \beta^{j-t} U(s_j, d_j, \varepsilon_j) | s_t, \varepsilon_t \right\} g(s_t|s) q(\varepsilon_t|s_t) d\varepsilon_t, \end{aligned} \tag{OS.1}$$

where the second equality follows from the definition of $V(s_t, \varepsilon_t)$ in (6). The expectation above is taken over the transition of (s_j, ε_j) starting from (s_t, ε_t) based on the transition probability $g(s'|s)$ and probability density $q(\varepsilon'|s')$. Thus, the function $\tilde{V}(s)$ represents the expected future utilities starting from intermediate state s and assuming the system always takes the optimal action.

We introduce following notations. We consider two systems s and s'. For every period t, we use s_t and s'_t to denote the system states at the start of period t, which are $s_t = \left(n^G_{l,t}, n^G_{h,t}, n^F_t\right)$ and $s'_t = \left(\left(n^G_{l,t}\right)', \left(n^G_{h,t}\right)', \left(n^F_t\right)'\right)$. Additionally, we use $\tilde{s}_t = \left(\tilde{n}^G_{l,t}, \tilde{n}^G_{h,t}, \tilde{n}^F_t\right)$ and $\tilde{s}'_t = \left(\left(\tilde{n}^G_{l,t}\right)', \left(\tilde{n}^G_{h,t}\right)', \left(\tilde{n}^F_t\right)'\right)$ to denote the intermediate states after actions d_t and d'_t are taken in period t, which are given by $\tilde{s}_t = \varphi(s_t, d_t)$ and $\tilde{s}'_t = \varphi(s'_t, d'_t)$ according to (8). Assume the two systems start from intermediate states $\tilde{s}_0$ and $\tilde{s}'_0$ with

$$\tilde{n}^G_{i,0} = \left(\tilde{n}^G_{i,0}\right)' \text{ for } i \in \{l, h\}, \text{ and } \tilde{n}^F_0 \leq \left(\tilde{n}^F_0\right)', \tag{OS.2}$$

then Proposition 2 translates to $\tilde{V}(\tilde{s}_0) \geq \tilde{V}(\tilde{s}'_0)$.

Coupling: Our proof is based on a coupling argument and induction in time. We first introduce the coupling of the two systems s and s' as follows: First, the two systems witness identical arrivals to GK and external arrivals to the FSU in every period, i.e., $A_{l,t} = A'_{l,t}$, $A_{h,t} = A'_{h,t}$, and $E_t = E'_t$ for every t. Next, we couple the FSU departures from the two systems as follows. Denote the numbers of FSU customers before departures by $\bar{n}_t^F$ and $(\bar{n}_t^F)'$ respectively and assume $\bar{n}_t^F \leq (\bar{n}_t^F)'$. Then, the departures D_t and D'_t are coupled as $D'_t = D_t + Z_t$, where D_t is a Binomial-$(\bar{n}_t^I, \mu_I)$ variable and Z_t is a Binomial-$((\bar{n}_t^F)' - \bar{n}_t^F, \mu_I)$ variable. That is, the number of departures in the s' system is always at least as many as the number in the s system. Finally, if an identical action d is taken in each system, the random utility components $\varepsilon_t(d)$ and $\varepsilon'_t(d)$ associated with that action d coincides for the two systems in every period t.

Under the coupling described above, we first prove following lemma.

Lemma 1 *Under the coupling, if the intermediate states in each system in period $t-1$ satisfy*

$$\tilde{n}_{i,t-1}^G = \left(\tilde{n}_{i,t-1}^G\right)' \textit{ for } i \in \{l,h\}, \textit{ and } \tilde{n}_{t-1}^F \leq \left(\tilde{n}_{t-1}^F\right)', \tag{OS.3}$$

then the states at the start of period t satisfy

$$n_{i,t}^G = \left(n_{i,t}^G\right)' \textit{ for } i \in \{l,h\} \textit{ and } n_t^F \leq \left(n_t^F\right)'. \tag{OS.4}$$

PROOF: The result follow directly from the coupled arrivals and departures in the two systems. Since we start from the intermediate state, the system evolution to period t is only dictated by the stochastic arrivals to the GK, external arrivals to the FSU, and departures from the FSU during period $t-1$.

It is trivial to see the relationship for the customers in GK holds as we assume same GK arrival processes by coupling. For the FSU customers, we consider both external arrivals and departures. We have

$$\begin{aligned}\left(n_t^F\right)' - n_t^F &= \min\left\{(\tilde{n}_{t-1}^F)' + E'_{t-1}, B\right\} - \min\left\{\tilde{n}_{t-1}^F + E_{t-1}, B\right\} - \left(D'_{t-1} - D_{t-1}\right) \\ &= \min\left\{(\tilde{n}_{t-1}^F)' + E_{t-1}, B\right\} - \min\left\{\tilde{n}_{t-1}^F + E_{t-1}, B\right\} - Z_{t-1} \geq 0.\end{aligned}$$

The last equality follows from the coupling $E_t = E'_t$ for external arrivals and $D_{t-1} + Z_{t-1} = D'_{t-1}$ for departures. The last inequality holds as Z_{t-1} follows a Binomial-(L_{t-1}, μ_I) distribution with $L_{t-1} = \min\left\{(\tilde{n}_{t-1}^F)' + E_{t-1}, B\right\} - \min\left\{\tilde{n}_{t-1}^F + E_{t-1}, B\right\}$. This completes the proof for Lemma 1. □

Mimicking Policy: We now define the policies used in each system. We assume the system s' always takes its optimal action which achieve the supremum in (OS.1). For system s, we define a mimicking policy π which mimics the action taken in the s' system whenever possible; if it is not possible, it takes its own optimal action. We denote the value function associated with this policy by $V^\pi(s)$, which is defined by (OS.1) with optimal action d_t replaced by the one under policy π. By definition, we have

$$\tilde{V}(\tilde{s}_0) \geq V^\pi(\tilde{s}_0). \tag{OS.5}$$

To prove the proposition, we will establish following two properties under our coupling and the policy π. First, two systems always have same number of customers in the GK, but system s has no more customers in the FSU:

$$n_{i,t}^G = \left(n_{i,t}^G\right)' \text{ for } i \in \{l,h\} \text{ and } n_t^F \leq \left(n_t^F\right)', \quad \forall t. \tag{OS.6}$$

Second, the action taken in the s' system is always admissible for system s; thus system s always mimics the action of s' under π:

$$d_t = d_t' \in \Pi(s_t), \quad \forall t. \tag{OS.7}$$

Note that (OS.6) directly implies (OS.7), as it follows from (1) that given the same number of customers in GK, the system with fewer FSU customers has a larger admissible action set, leading to $d_t' \in \Pi(s_t') \subseteq \Pi(s_t)$.

Induction: We establish (OS.6) for every t by induction.

Base Case: The base case follows directly from the relationship of the initial intermediate states $\tilde{s}_0$ and $\tilde{s}_0'$, which satisfy (OS.2), and from Lemma 1. Thus we have $n_{i,1}^G = \left(n_{i,1}^G\right)'$ for $i \in \{l,h\}$ and $n_1^F \leq \left(n_1^F\right)'$.

Inductive Step: We assume (OS.6) holds for period j and show this implies it holds for period $j+1$. In period j, under policy π, system s takes the same action of s' since by the inductive hypothesis the action is admissible, i.e., $d_j = d_j'$. Given the same action is taken in each system, it is easy to verify the intermediate states after action also coincide, i.e., $\tilde{n}_{i,j}^G = \left(\tilde{n}_{i,j}^G\right)'$ for $i \in \{l,h\}$ and $\tilde{n}_j^F = \left(\tilde{n}_j^F\right)'$. Finally, we can apply Lemma 1 to prove the relationship (OS.6) holds for period $j+1$. This completes the inductive step.

Per-Period Utilities: We have shown that under our coupling and the policy π for the s system, at the start of each period, the two systems always have same numbers of customers in the GK, and system s always has fewer customers in the FSU than that in system s'. Thus, the system s always mimics the action by s' under the policy π. We next prove the per-period utilities always coincide for the two systems, which follows by:

$$\begin{aligned} U(s_t, d_t, \varepsilon_t) &= u(s_t, d_t) + \varepsilon_t(d_t) = \sum_{i \in \{l,h\}} u_{a,i} a_{i,t} + \sum_{i \in \{l,h\}} u_{r,i} r_{i,t} + \sum_{i \in \{l,h\}} u_{w,i} \left(n_{i,t}^G - a_{i,t} - r_{i,t}\right) + \varepsilon_t(d_t) \\ &= \sum_{i \in \{l,h\}} u_{a,i} a_{i,t}' + \sum_{i \in \{l,h\}} u_{r,i} r_{i,t}' + \sum_{i \in \{l,h\}} u_{w,i} \left(\left(n_{i,t}^G\right)' - a_{i,t}' - r_{i,t}'\right) + \varepsilon_t'(d_t') = U(s_t', d_t', \varepsilon_t'). \end{aligned}$$

This is because: (i) Both systems take the same action, thus they send same numbers of customers to the FSU and SSU, leading to same routing utilities; (ii) As both systems have same numbers of customers in the GK, the number of customers remaining in the GK after actions are also the same, leading to the same waiting utilities; (iii) By our coupling, the random utility components coincide for the same action $d_t = d_t'$, i.e., $\varepsilon_t(d_t) = \varepsilon_t'(d_t')$.

As the per-period utilities coincide for every period given system s takes policy π and system s' takes its own optimal policy, we have $V^\pi(\tilde{s}_0) = \tilde{V}(\tilde{s}_0')$. Then it follows by (OS.5)

$$\tilde{V}(\tilde{s}_0) \geq V^\pi(\tilde{s}_0) = \tilde{V}(\tilde{s}_0').$$

This proves the proposition.

OS.1.2. Proof for Proposition 3

We prove equation (11) in below. The proof of equation (12) follow in a similar way. For the two states $s_1 = (1,0,0)$ and $s_2 = (0,1,0)$, denote the probabilities of admitting the patient to FSU and SSU under the them by $\Pr(a|s_1)$, $\Pr(r|s_2)$, $\Pr(a|s_1)$, and $\Pr(r|s_2)$ respectively. By Proposition 1, we can compute the log of these choice probabilities as

$$\ln \Pr(a|s_1) = \beta \tilde{V}(s_a) - C(s_1), \quad \ln \Pr(r|s_1) = -c_{r,l} + \beta \tilde{V}(s_r) - C(s_1), \tag{OS.8}$$

and

$$\ln \Pr(a|s_2) = \beta \tilde{V}(s_a) - C(s_2), \quad \ln \Pr(r|s_2) = -c_{r,h} + \beta \tilde{V}(s_r) - \ C(s_2), \tag{OS.9}$$

where $s_a = (0,0,1)$ and $s_r = (0,0,0)$ denote the system state after the customer is admitted to the FSU and SSU respectively. Here we use the assumption $u_{a,l} = u_{a,h} = 0$. Taking the differences in the choice probabilities across the two states for FSU and SSU admission decisions respectively, we obtain

$$\ln \Pr(a|s_1) - \ln \Pr(a|s_2) = C(s_2) - C(s_1). \tag{OS.10a}$$

and

$$\ln \Pr(r|s_1) - \ln \Pr(r|s_2) = u_{r,l} - u_{r,h} + C(s_2) - C(s_1). \tag{OS.10b}$$

Here we see the terms involving the value function $\tilde{V}(\cdot)$ in (OS.8) – (OS.9) are canceled out in the differences. Further subtracting (OS.10a) from (OS.10b), we get rid of the terms related to state-dependent function $C(\cdot)$ and identify the SSU admission utility difference as

$$u_{r,l} - u_{r,h} = \ln\left(\frac{\Pr(r|s_1)}{\Pr(a|s_1)}\right) - \ln\left(\frac{\Pr(r|s_2)}{\Pr(a|s_2)}\right).$$

This proves (11). The above steps show how well-constructed state and action pairs can be used to disentangle the complex structure in the choice probability expression (9). The proof of (12) follows similarly.

OS.1.3. Proof for Proposition 4

With non-zero utility $u_{a,h}$ for sending the high-priority customer to the FSU, the per-period utility of action is given by

$$\begin{aligned} u'(s_t, d_t) &= u_{a,l} a_{l,t} + u_{r,l} r_{l,t} + u_{w,l}\left(n^G_{l,t} - a_{l,t} - r_{l,t}\right) + u_{a,h} a_{h,t} + u_{r,h} r_{h,t} + u_{w,h}\left(n^G_{h,t} - a_{h,t} - r_{h,t}\right) \\ &= \sum_{\iota \in \{l,h\}} \Delta u_{w,\iota}\left(n^G_{\iota,t} - a_{\iota,t} - r_{\iota,t}\right) + \sum_{\iota \in \{l,h\}} \Delta u_{r,\iota} r_{\iota,t} + \Delta u_{a,l} a_{l,t} + u_{a,h} n^G_t, \end{aligned} \tag{OS.11}$$

where $n^G_t = n^G_{l,t} + n^G_{h,t}$ denotes the total number of customers in GK in period t. In addition, we denote $\Delta u_{w,\iota} = u_{w,\iota} - u_{a,h}$, and $\Delta u_{r,\iota} = u_{r,\iota} - u_{a,h}$ for $\iota \in \{l,h\}$, as well as $\Delta u_{a,l} = u_{a,l} - u_{a,h}$. They represent the relative actions utilities compared with sending the high-priority customer to the FSU.

We see that, with non-zero utility for sending high-priority customer to the FSU, the last term $u_{a,h}n_t^G$ is involved in the per-period utility. Denote the per-period utility and value function with non-zero FSU routing utility $u_{a,h}$ by $U'(s_t,d_t,\varepsilon_t)$ and $V'(s_t,\varepsilon_t)$, respectively. By (6), we have

$$V'(s_t,\varepsilon_t)=\sup_{d_t\in\Pi(s_t)}\mathsf{E}\left\{\sum_{j=t}^{\infty}\beta^{j-t}U'(s_j,d_j,\varepsilon_j)|s_t,\varepsilon_t\right\}=\sup_{d_t\in\Pi(s_t)}\mathsf{E}\left\{\sum_{j=t}^{\infty}\beta^{j-t}\left[U(s_j,d_j,\varepsilon_j)+u_{a,h}n_j^G\right]|s_t,\varepsilon_t\right\}, \tag{OS.12}$$

where $U(s_j,d_j,\varepsilon_j)$ denotes the per-period utility with zero FSU routing utility $u_{a,h}$ and same relative utilities of other actions, i.e., $\Delta u_{w,\iota}$ and $\Delta u_{r,\iota}$ for $\iota\in\{l,h\}$ as well as $\Delta u_{a,l}$.

With same total number of customers in GK for each period, we can drop the subscript j in n_j^G and write it as n^G. Then the term $u_{a,h}n^G$ in (OS.12) is independent of action d_t and can be extracted out of the conditional expectation. We can write (OS.12) as

$$V'(s_t,\varepsilon_t)=\sup_{d_t\in\Pi(s_t)}\mathsf{E}\left\{\sum_{j=t}^{\infty}\beta^{j-t}U(s_j,d_j,\varepsilon_j)|s_t,\varepsilon_t\right\}+\frac{1}{1-\beta}u_{a,h}n^G, \tag{OS.13}$$

where the last term is a constant number and independent of the optimal actions. We note that by (6), the first term in (OS.13) is exactly the value function under the model with zero FSU routing utility for high-priority customers and same relative utilities for other actions. Thus, we can establish the following relationship between the value functions with different normalization levels of $u_{a,h}$:

$$V'(s_t,\varepsilon_t)=V(s_t,\varepsilon_t)+\frac{1}{1-\beta}u_{a,h}n^G.$$

Similarly, by (OS.1), we can prove the function $\tilde{V}'(s)$ under non-zero utility $u_{a,h}$ satisfies

$$\tilde{V}'(s)=\tilde{V}(s)+\frac{1}{1-\beta}u_{a,h}n^G. \tag{OS.14}$$

By (OS.11) and (OS.14), we have

$$u'(s,d)+\beta\tilde{V}'(\varphi(s,d))=u(s,d)+\beta\tilde{V}(\varphi(s,d))+\frac{1}{1-\beta}u_{a,h}n^G.$$

Plugging this in to (10), it is easy to verify the new function $\tilde{V}'(s)$ satisfies the functional equation. Moreover, by (9), we can show that the new choice probability $f'(d_t|s_t)$ coincides with the original ones under $u_{a,h}=0$. Thus, model identification is not affected by the normalization level of $u_{a,h}$.

OS.2. Goodness of Fit of Estimated Structural Model

We show our estimated structural model provides good fits for both the hospitals' decisions and the overall system statistics. First, we compare the explanatory power for the hospitals' decisions, as measured by the McFadden's pseudo R^2, from our structural model with that from the reduced-form multinomial logistic regression model in Section EC.3 of the Online Supplement. As shown in Tables 1 and EC.1, the pseudo

R^2 for the combined estimation of all hospitals is similar from the structural model and logistic regression (0.14 versus 0.16). The average pseudo R^2 for individual hospital estimations is also comparable (0.13 vs 0.16). We note that the multinomial model contains a comprehensive set of variables that might influence admission decisions, including patient's characteristics, system states, and multiple dummies for seasonality fixed effects. Thus, the similar R^2 suggests that our structural model has reasonable explanatory power in capturing the hospitals' decisions.

We also find the structural model and reduced-form regression provide consistent evidences for the heterogeneity in hospitals' discounting behaviors. In the multinomial logistic model, we measure the impact of ICU occupancy on ICU admission likelihood by the coefficient γ_{ICU}. As discussed in Section EC.3, γ_{ICU} is significantly negative for most hospitals, suggesting the ICU admission probability drops as ICU occupancy increases. We find that the estimated $\hat{\beta}$ and γ_{ICU} are negatively correlated: the correlation is -0.770 across the 22 hospitals, statistically significant at the 0.1% level. As such, hospitals with larger estimated discount factors indeed tend to be more responsive to ICU congestion when admitting patients.

Next, we use simulation to show that our structural model produces system statistics close to those observed in the data. We examine several important system statistics, including the average ICU occupancy, overall proportion eventually admitted to the ICU for the two classes, as well as the ED waiting times for the two classes. The statistics estimated from our structural model are averaged from 200 simulation trials as in the counterfactual studies in Section 5.

Figure OS.1 Comparison of System Statistics from Structural Model and Real Data

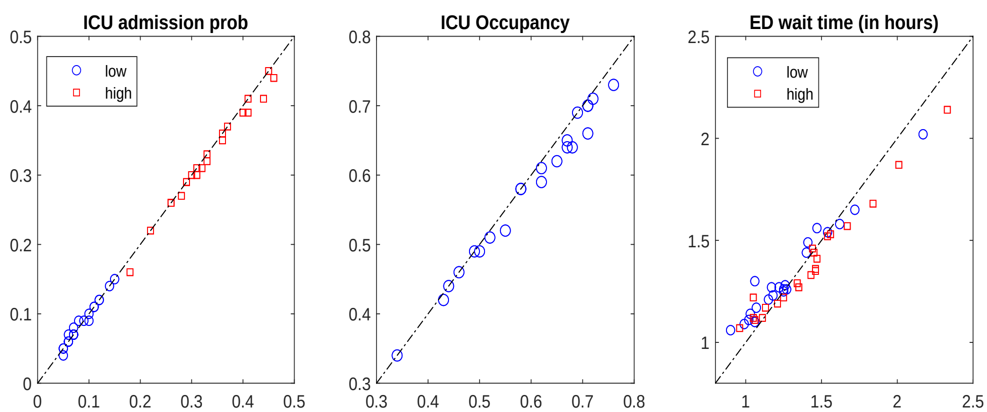

Note: The figure compares the system statistics simulated from the structural model (y-coordinate) and observed from the real data (x-coordinate), including ICU admission probability (left), ICU occupancy (middle), and ED waiting times (right).

Figure OS.1 compares the simulated and observed system statistics: ICU admission probabilities (left), ICU occupancy (middle), and ED waiting times (right). In the panels, each point represents a hospital in our study; its x-coordinate (y-coordinate) corresponds to the observed (simulated) value of the system statistic.

We plot the 45-degree line in each panel, which represents a perfect fit. As we can see, most points fall very close to the 45-degree line, implying that our estimated structural model produces system statistics that are very close to the observed data. We also check the correlation and the relative mean absolute deviation (MAD) between the observed and simulated system statistics. The results are reported in Table OS.1. We see the correlations are very close to one and the relative deviation is small (2.2% to 4.5%). This further supports its effectiveness in modeling the admission process for ED patients.

Table OS.1 Correlation and Relative MAD between Simulated and Observed Statistics

	ICU Adm. Prob	ICU Avg Occu	Avg ED Wait
Correlation	0.999	0.993	0.976
Relative MAD	2.21%	2.58%	4.49%

With a small number of parameters, over-fitting is unlikely for our structural model. To further address this concern, we divide the sample to first and second halves for each hospital. We estimate the structural model with all hospitals combined using the first half sample, and do out-of-sample prediction on the second half. The McFadden's pseudo R^2 from the out-of-sample prediction (0.13) is very similar to that from the in-sample estimation (0.14) and the level from the full sample estimation in Table 1 (0.14).

OS.3. Impact of ICU Occupancy on Patient Outcomes

In this section, we use reduced-form analysis to examine if ICU congestion affects the outcomes of patients admitted to the ICU. We run an OLS regression to examine how ICU occupancy affects patient outcomes:

$$Y_i = \gamma_{0,d} + \gamma_{L,d}\text{LAPS2}_i + \gamma_{ICU,d}\text{ICUOccu}_t + \gamma_{ED,d}\text{EDNum}_t + \boldsymbol{\gamma}_{Z,d}^\top \mathbf{Z}_{i,t} + \epsilon_{it}, \tag{OS.15}$$

where Y_i is one of the following patient outcomes: 30-day mortality, hospital readmission, and hospital length-of-stay (LOS). ICUOccu_t is the ICU occupancy level at the time of patient admission, measured by the percentage occupancy relative to the ICU capacity; EDNum_t is the number of patients in the ED (including both high and low severity types). The other patient and system-level controls $\mathbf{Z}_{i,t}$ are introduced in Section EC.3. The set-up of the OLS regression (OS.15) is similar to our reduced-form analysis for the patient admission decision in Section EC.3.

As an additional test, we also examine the impact of the average ICU occupancy during a patient's ICU stay, rather than the occupancy level at the time of ICU admission. The average ICU occupancy is also measured relative to the ICU capacity.

$$Y_i = \gamma_{0,d} + \gamma_{L,d}\text{LAPS2}_i + \gamma_{ICU,d}\text{AvgICUOccu}_i + \gamma_{ED,d}\text{EDNum}_t + \boldsymbol{\gamma}_{Z,d}^\top \mathbf{Z}_{i,t} + \epsilon_{it}, \tag{OS.16}$$

The average ICU occupancy, AvgICUOccu_i, is computed using the two-hour snapshots from our structural model. For example, if a patient enters the ICU at 12:45 and leaves at 17:16, we take the average of ICU occupancy levels recorded at the 13:00, 15:00, and 17:00 snapshots. Note the average ICU occupancy level is not available to the physician at the time of a patient's ICU admission, but this analysis allows us to understand the association between congestion and outcomes. We replace ICUOccu_t by AvgICUOccu_i in regression (OS.15). The two variables are not included simultaneously as they are highly correlated (a correlation of 0.85).

Table OS.2 Estimation Results for OLS Regression (OS.15), N = 19,680

Panel A: ICU occupancy at the time of admission

	Mortality	Readmission	Hospital LOS
LAPS2	0.004^{***}	0.0006^{***}	0.0355^{***}
	(8.02×10^{-5})	(8.68×10^{-5})	(0.002)
ICUOccu_t	-0.0076	0.0022	-0.381
	(0.017)	(0.019)	(0.491)
EDNum_t	-0.0009	0.0048^{*}	0.0898
	(0.002)	(0.002)	(0.055)

Panel B: Average ICU occupancy during a patient's stay

	Mortality	Readmission	Hospital LOS
LAPS2	0.0004^{***}	0.0006^{***}	0.0354^{***}
	(8.01×10^{-5})	(8.68×10^{-5})	(0.002)
AvgICUOccu_i	-0.011	0.0153	-0.0423
	(0.020)	(0.021)	(0.556)
EDNum_t	-0.0008	0.0047^{*}	$0.0897^{\dagger}$
	(0.002)	(0.002)	(0.053)

Note: $^{\dagger}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, and $^{***}p < 0.001$.

Since we focus on the ED medical patients in our structural model, we estimate (OS.15) and (OS.16) using the sample of ED medical patients admitted to the ICU. The estimation results are reported in Table OS.2. The coefficients of ICU occupancy level are statistically insignificant for all three outcomes. This holds for both the ICU occupancy level at the time of admission (Panel A) and the average ICU occupancy level during a patient's stay (Panel B). We also note that the magnitudes of the estimated coefficients are small across all three outcomes and two specifications, except for LOS in Panel A, where the impact of ICU occupancy on hospital LOS is negative–i.e., higher ICU occupancy is possibly associated with shorter hospital LOS. These results suggest that, even if we had larger sample size, it is still unlikely that we would find higher ICU occupancy associated with worse patient outcomes. Therefore, we do not find empirical evidence that ICU congestion leads to worse patient outcomes in our study cohort.

As expected, the LAPS2 severity score is significantly positive for all three outcomes. In addition, we find that the number of patients in the ED at the time of ICU admission is associated with a higher readmission rate and a longer hospital length of stay. The latter effect is statistically significant only at the 10% level, and only when we control for the average ICU occupancy during a patient's stay (Panel B). One possible explanation for the ED occupancy effect is that higher ED occupancy may delay the initiation of care due to longer boarding time. In addition, the care quality in the ED may be more sensitive to patient occupancy than that in the ICU, as the ICU typically operates under a fixed nurse-to-bed ratio.

OS.4. Dynamic Model with Intraday Time Trend

In this section, we extend our dynamic discrete choice model for the ICU admission problem by including a time indicator in the state. This captures the time trend in the system transition during a day, e.g., the arrival and departure rates may vary substantially during day time and night time.

The model is set up as follows. Recall our data consists of state-action pairs for each two-hour interval. We now include a time indicator $m_t \in \{1, 2, ..., 12\}$ in the state variable, which denotes the specific time (two-hour interval) of the day. For example, $m_t = 1$ means the time of the day is from 12AM to 2AM. Thus, our new state variable becomes $s_t = (m_t, n^G_{l,t}, n^G_{h,t}, n^F_t)$. Naturally, the transition of m_t is given by

$$m_{t+1} = \begin{cases} m_t + 1, & \text{if } m_t < 12 \\ 1, & \text{if } m_t = 12. \end{cases}$$

To capture the time trend in the system state, we estimate the ED arrival rates ($\lambda_{Q,l}$ and $\lambda_{Q,h}$), external ICU arrival rate (λ_E), and ICU departure rate (μ_I) for each of the 12 two-hour time intervals separately. That is, we now estimate $\{\lambda^{(m)}_{Q,l}\}$, $\{\lambda^{(m)}_{Q,h}\}$, and $\{\lambda^{(m)}_{E}\}$ for $m = 1, 2, ..., 12$ following the procedures in Section EC.2.2. This accounts for the intra-day variation in the system state transition. In other words, we allow the state transition probability $g(s'|s)$ of $(n^G_{l,t}, n^G_{h,t}, n^F_t)$ derived in Section EC.1.1 to depend on the time interval m_t using the estimated arrival and departure rates for the specific time period. The ED and ICU capacities are still fixed as they do not vary over time.

We estimate the new model using our data. However, including the time indicator increases our state space by 12 times. To reduce computational burden, we now drop the buffer term in the ED capacity (see Section EC.2.2) and set it for each severity class as the maximum number of patients observed in the data. This has limited impact on the estimation since, in most cases, there are usually only one or two patients in the ED. Like our main specification, we consider a discrete grid $\{0.1, 0.2, ..., 0.9\}$ for potential discount factor levels.

The estimation results for the time-trend model are reported in Table OS.3 for all hospitals combined and Table OS.12 for individual hospitals. The $\hat{\beta}$ column reports the estimated discount factor, and the four subsequent columns ($\hat{u}_{w,l}$, $\hat{u}_{r,l}$, $\hat{u}_{w,h}$, $\hat{u}_{r,h}$) report the estimated utility parameters for the low and high severity classes. By Table OS.3, we see that the estimated model parameters under the time-trend model are very similar to those in our main specification. In particular, the estimated discount factor $\hat{\beta}$ moves from 0.3 to the adjacent level of 0.2. The utility estimates are also very similar.

Table OS.3 Comparison of Main Model and Time-trend Model: All Hospitals Combined

	$\hat{\beta}$	$\hat{u}_{w,l}$	$\hat{u}_{r,l}$	$\hat{u}_{w,h}$	$\hat{u}_{r,h}$
Main specification	0.3	-0.071	1.950	-0.932	0.674
Time-trend model	0.2	-0.003	1.887	-0.839	0.680

In Table OS.12 for individual hospitals, the last "$\hat{\beta}$ (main)" column additionally reports the estimated discount factor in our main setting without time trend. We see that the estimated discount factors are very similar for most hospitals between the two specifications. In particular, the estimated discount factors are exactly the same for 15 out of the 22 hospitals and changes to adjacent levels for 5 hospitals. Only for two hospitals (Hospital 7 and 11), the estimated $\hat{\beta}$ experiences a larger change after the time trend is incorporated. For these hospitals, we have a smaller sample size and thus less variation. This explains the change in the discount factor once we incorporate the time trend. The above results suggest that our main findings are robust to including the time trend in the dynamic discrete choice model.

OS.5. Dynamic Model with Three Patient Classes

To further capture the heterogeneity in patient's utility coefficients, we estimate a dynamic structure model with three patient classes. The three severity classes include patients with severity scores (LAPS2 score) below the 60th percentile, between the 60th and 85th percentile, and those above the 85th percentile, respectively. We normalize the ICU admission utility to zero for all three classes and allow them to have different utility parameters for ED waiting and non-ICU admission. The arrival rates and ED capacity of the three classes are estimated directly from data as in our main set-up. Thus, we have in total seven parameters to estimate in the identification, including the discount factor, and the waiting and non-ICU admission utilities for the three classes. The discount factor is still estimated on the grid $\{0.1, 0.2, ..., 0.9\}$. To control the state space dimension, we drop the buffer term in the ED capacity and set it for each class as the maximum number of patients in the data.

The estimation results for three-class model are reported in Table OS.4 for all hospitals combined and Table OS.13 for individual hospitals. By Table OS.4, we see that our estimation results remain very robust when using the three-class model. The identified discount factor is still 0.3. The estimated utility parameters are also similar. Recall that the high severity class is defined for the same patient cohort in the main model and the three-class model (the patients with LAPS2 above the 85th percentile). We see that their utility estimates $u_{w,h}$ and $u_{r,h}$ are indeed very close in the two models (-0.972 vs. -0.950 for $u_{w,l}$ and 0.674 vs 0.684 for $u_{r,h}$).

Table OS.4 Comparison of Main Model and Three-class Model: All Hospitals Combined

Main specification					Three-class model						
$\hat{\beta}$	$\hat{u}_{w,l}$	$\hat{u}_{r,l}$	$\hat{u}_{w,h}$	$\hat{u}_{r,h}$	$\hat{\beta}$	$\hat{u}_{w,l}$	$\hat{u}_{r,l}$	$\hat{u}_{w,m}$	$\hat{u}_{r,m}$	$\hat{u}_{w,h}$	$\hat{u}_{r,h}$
0.3	-0.071	1.95	-0.932	0.674	0.3	-0.001	2.299	-0.418	1.794	-0.95	0.684

Table OS.13 reports the estimation results with three patient classes for individual hospitals. The second column reports the dimensions of state space in our three-class model, which are indeed relatively large.

The third column ($\hat{\beta}_{3c}$) reports the estimated discount factor in the three-class model. The subsequent three columns ($\hat{u}_{w,1}$, $\hat{u}_{w,2}$, and $\hat{u}_{w,3}$) report the estimated ED waiting utilities for the low-, medium-, and high-severity patients. The next three columns ($\hat{u}_{r,1}$, $\hat{u}_{r,2}$, and $\hat{u}_{r,3}$) report the estimated non-ICU admission utilities for the three classes. In the last column, we report the estimated discount factor from the two-class model as in our main manuscript.

We have the following observations about the results. First, the estimated discount factors remain similar to the main specification with two patient classes for most of the individual hospitals. The absolute difference between $\hat{\beta}_{2c}$ and $\hat{\beta}_{3c}$ is zero for 13 hospitals, and 0.1 for 8 hospitals This suggests that our main identification and estimation results are robust under the three-class model. Second, we find that for most hospitals, the estimated ED waiting and non-ICU admission utilities are also similar to those in the main specification and are decreasing in patient severity level. This also makes sense as more severe patients on average benefit more from the ICU admission (recall the ICU admission utilities are normalized to zero for all three classes). It also suggests that our main model with two classes of patients captures most of the patient heterogeneity in the sample.

OS.6. Supplementary Tables and Figures

Figure OS.2 Probability of ICU admission by ICU occupancy rates (selected hospitals)

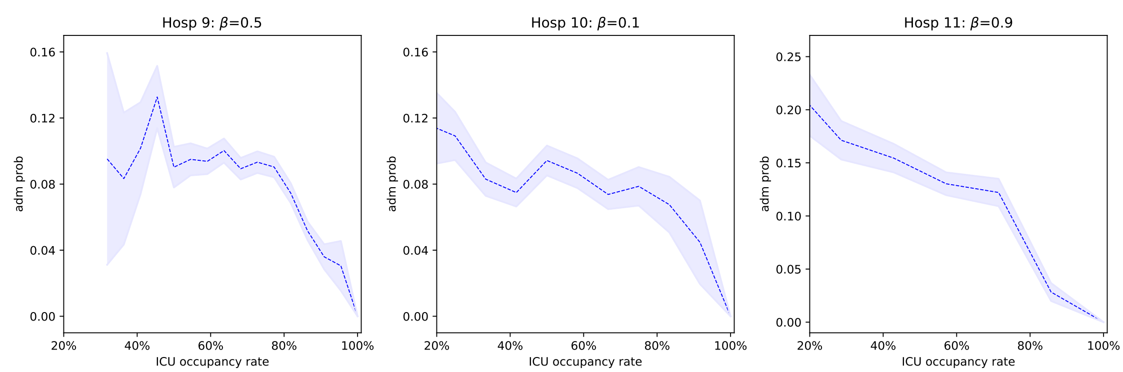

Table OS.5 Summary Statistics of ED Medical Patient Characteristics

Final study cohort: N=164,167

	Min	Max	Mean	Median	SD
Male			0.53		
Age (years)	18.00	113.00	67.27	70.00	17.59
LAPS2	0.00	294.00	74.11	70.00	37.47
EDWait (hours)	0.02	12.00	1.30	0.88	1.41

ICU admission cohort: N=19,683

	Min	Max	Mean	Median	SD
Male			0.48		
Age	18.00	111.00	64.52	67.00	17.48
LAPS2	0.00	294.00	105.03	102.00	45.98
EDWait	0.02	11.98	1.36	0.90	1.45

Note: LAPS2 is severity of illness score. EDWait corresponds to the ED boarding time

Figure OS.3 Utility Estimates for Each β' in the Counterexample

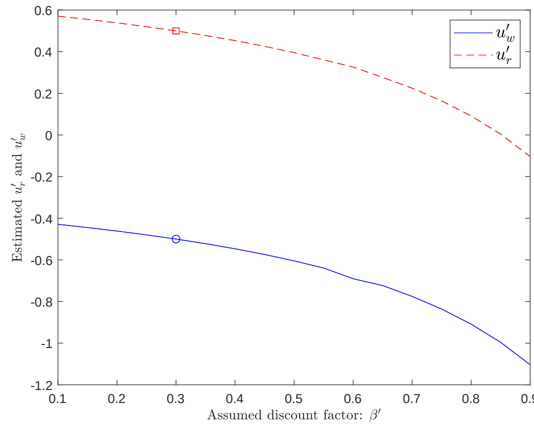

Figure OS.4 Log Likelihood at Different β for Select ED States (All Hospitals Combined)

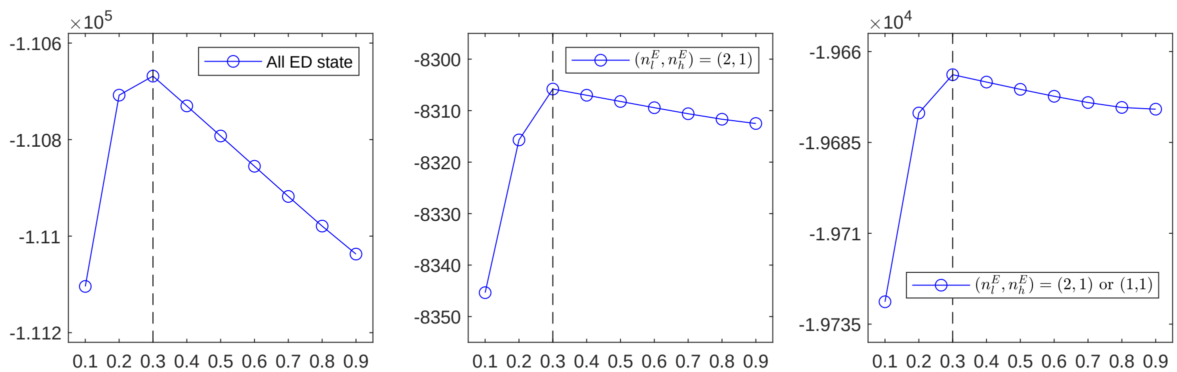

Figure OS.5 Examples of Log-likelihood versus Discount Factor for a Subset of Hospitals

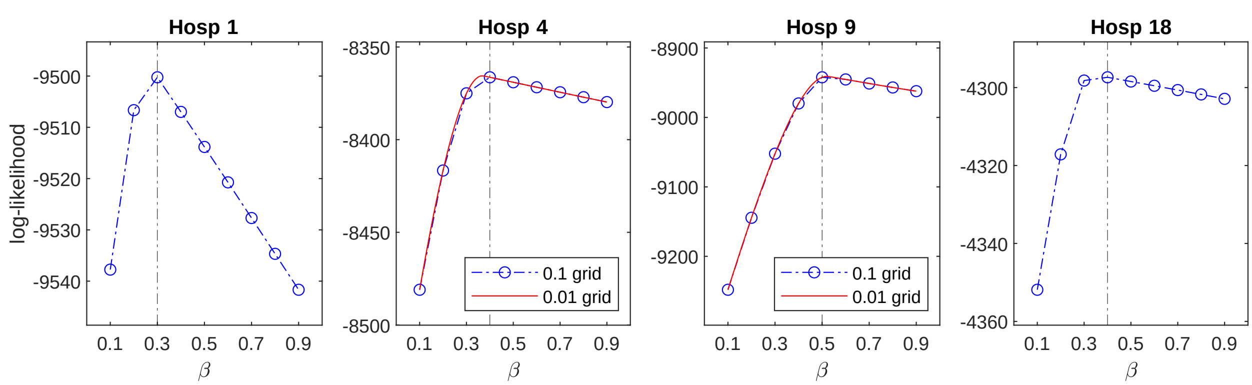

Note: The estimated log-likelihood at $\beta = \{0.1, 0.2, \ldots, 0.9\}$ for Hospitals 1, 4, 9, and 18. The β with the best likelihood is plotted by the black vertical line. The red solid line shows the likelihood with a 0.01 grid for Hospitals 4 and 9.

Table OS.6 System Summary Statistics by Hospital

Hosp	Num. of days	Num. of hospitalizations	Q_l	Q_h	B	ICUOccu	$\lambda_{Q,l}$	$\lambda_{Q,h}$	λ_E	μ_I	$\Pr(a_l)$	$\Pr(a_h)$
1	667	11,676	11	5	21	0.67	1.230	0.221	0.252	0.035	0.12	0.41
2	500	9,902	14	7	26	0.76	1.429	0.227	0.268	0.026	0.11	0.36
3	667	8,039	8	4	12	0.49	0.859	0.146	0.101	0.030	0.05	0.22
4	667	14,595	13	7	31	0.71	1.529	0.327	0.519	0.031	0.05	0.29
5	576	5,082	7	3	11	0.65	0.633	0.098	0.108	0.030	0.11	0.41
6	667	10,577	9	5	21	0.58	1.123	0.208	0.278	0.036	0.08	0.36
7	578	4,915	7	4	11	0.71	0.583	0.111	0.188	0.030	0.05	0.18
8	653	8,400	10	6	16	0.67	0.868	0.185	0.158	0.033	0.14	0.40
9	514	12,355	15	5	22	0.71	1.807	0.264	0.354	0.036	0.07	0.31
10	609	5,978	8	4	12	0.52	0.714	0.092	0.228	0.048	0.06	0.28
11	667	2,655	6	4	7	0.55	0.284	0.045	0.056	0.029	0.12	0.46
12	388	6,751	12	6	24	0.69	1.256	0.173	0.299	0.026	0.06	0.33
13	667	8,061	9	4	16	0.50	0.817	0.182	0.110	0.028	0.07	0.30
14	667	12,841	11	6	36	0.72	1.332	0.257	0.493	0.026	0.07	0.33
15	575	7,208	8	4	16	0.43	0.897	0.111	0.114	0.031	0.07	0.31
16	667	7,190	8	4	13	0.44	0.752	0.133	0.108	0.033	0.06	0.26
17	548	3,511	7	4	9	0.62	0.418	0.099	0.066	0.024	0.09	0.32
18	667	7,702	9	6	32	0.58	0.796	0.122	0.557	0.034	0.06	0.26
19	547	7,476	8	4	25	0.46	0.966	0.170	0.196	0.031	0.10	0.37
20	666	5,096	8	4	11	0.34	0.545	0.087	0.065	0.035	0.07	0.29
21	333	2,109	8	4	13	0.68	0.451	0.076	0.139	0.025	0.10	0.44
22	333	2,048	8	4	16	0.62	0.422	0.090	0.182	0.028	0.15	0.45

Note: System summary statistics for each hospital: Q_i for $i \in \{l, h\}$ is the ED capacity for the two classes of patients; B is the ICU capacity; ICUOccu is the average ICU occupancy level; $\lambda_{Q,i}$ for $i \in \{l, h\}$ is the ED arrival rate; λ_E is the external arrival rate to ICU; μ_I is the ICU departure rate; $\Pr(a_i)$ for $i \in \{l, h\}$ is the overall ICU admission probability for the ED patients.

Table OS.7 Hospital Characteristics Considered in the Heterogeneity Analysis

Size	ICU size, ED size for low and high severity classes
Arrival rates	ED arrival rate for low and high severity classes, external ICU arrival rate
Departure rate	ICU departure rate
ICU congestion	Average ICU occupancy (as % of capacity), median ICU occupancy (as % of capacity), % of periods with fewer than one or two empty ICU beds
Variation in system	Standard deviations of ED arrivals, ICU external arrivals, and ICU occupancy across periods
Patient characteristics	average age, average LAPS2 score, average ED waiting time for two classes, % of high severity class patients, % of ICU admissions from external arrivals
Utility parameters	waiting and non-ICU admission utilities for both classes

Table OS.8 Counterfactual Estimates of Impact when β Increases from $\hat{\beta}$ to 0.9 (All Hospitals)

Hosp	$\hat{\beta}$	ΔDays HighCgstn (in days)			ΔPats HighCgstn (in patients)			ΔPats Balk (in patients)			Increase in ED Time (in hours)	
		β to 0.9	$B+1$	Heuris	β to 0.9	$B+1$	Heuris	β to 0.9	$B+1$	Heuris	Low Class	High Class
1	0.3	4.19	6.27	3.38	34.94	52.26	28.20	4.68	6.73	4.27	2.05	0.63
2	0.5	6.04	6.78	6.44	47.03	52.78	50.12	7.53	8.43	8.73	1.84	0.63
3	0.1	1.75	6.06	2.26	6.89	23.82	8.88	0.81	2.11	0.77	1.59	0.29
4	0.4	1.60	4.22	1.17	17.89	47.08	13.11	3.95	8.63	2.81	1.33	0.42
5	0.1	6.79	18.74	5.24	24.67	68.04	19.02	4.55	8.84	3.11	1.75	0.44
6	0.2	1.15	2.73	1.36	10.07	23.77	11.83	1.27	2.42	1.52	1.70	0.50
7	0.9	-0.45	23.23	0.01	-1.62	84.04	0.04	-0.71	21.20	1.33	0.00	0.00
8	0.1	5.58	10.90	5.45	33.21	64.88	32.41	4.69	7.92	4.62	1.36	0.76
9	0.5	4.27	8.60	6.37	38.99	78.61	58.22	6.42	13.57	10.56	1.55	0.49
10	0.1	0.81	8.23	1.09	5.10	51.80	6.87	0.81	7.56	1.51	0.73	0.26
11	0.9	-0.11	26.98	2.79	-0.24	56.65	5.86	-0.04	6.61	0.54	0.00	0.00
12	0.6	2.14	6.06	2.96	15.20	43.17	21.05	3.35	7.67	4.07	1.15	0.71
13	0.1	0.96	2.34	0.79	4.79	11.65	3.92	0.50	0.95	0.32	2.08	0.38
14	0.3	1.58	2.36	1.73	17.43	25.90	19.07	3.26	5.26	3.54	1.26	0.46
15	0.1	0.22	0.85	0.11	1.25	4.77	0.64	0.13	0.30	0.00	0.37	0.11
16	0.1	0.65	2.52	0.95	3.09	12.07	4.56	0.29	0.95	0.52	1.48	0.44
17	0.9	0.35	21.54	0.65	0.82	50.58	1.52	-0.03	6.60	0.35	0.00	0.00
18	0.4	0.08	0.65	0.16	1.07	8.24	2.05	0.23	1.16	0.62	1.42	0.28
19	0.1	0.05	0.08	0.02	0.40	0.75	0.20	0.04	0.09	0.05	0.45	0.16
20	0.3	0.25	0.94	0.06	1.05	3.91	0.26	0.02	0.15	0.08	1.90	0.33
21	0.7	1.90	16.61	2.71	6.77	59.02	9.62	1.86	11.97	2.59	1.33	0.49
22	0.3	1.12	8.27	1.35	9.75	42.31	6.91	2.43	7.59	2.29	3.23	0.47

Table OS.9 **Comparison of Identified Discount Factor $\hat{\beta}$ and Coefficient of ICU Occupancy γ_{ICU}**

Hosp	Discount factor $\hat{\beta}$ Structural model	Coefficient γ_{ICU} Full logit model	Coefficient γ_{ICU} Simplified logit model
1	0.3	-0.372^{**}	$-0.162^{\dagger}$
	(0.012)	(0.114)	(0.089)
2	0.5	-0.642^{***}	-0.445^{***}
	(0.010)	(0.116)	(0.099)
3	0.1	-0.191	-0.175
	(0.019)	(0.193)	(0.153)
4	0.4	-0.344^{*}	-0.020
	(0.010)	(0.140)	(0.105)
5	0.1	-0.782^{***}	-0.488^{***}
	(0.023)	(0.183)	(0.140)
6	0.2	$-0.270^{\dagger}$	0.144
	(0.014)	(0.146)	(0.107)
7	0.9	-1.489^{***}	-1.170^{***}
	(0.016)	(0.230)	(0.208)
8	0.1	-0.591^{***}	-0.404^{***}
	(0.020)	(0.117)	(0.099)
9	0.5	-0.419^{***}	-0.258^{**}
	(0.009)	(0.126)	(0.104)
10	0.1	-0.434	$-0.275^{\dagger}$
	(0.027)	(0.201)	(0.166)
11	0.9	-0.969^{***}	-0.956^{***}
	(0.017)	(0.216)	(0.185)
12	0.6	-1.025^{***}	-0.508^{***}
	(0.010)	(0.219)	(0.149)
13	0.1	-0.214	-0.06
	(0.018)	(0.184)	(0.131)
14	0.3	-0.269^{*}	-0.127
	(0.012)	(0.132)	(0.100)
15	0.1	-0.083	-0.035
	(0.033)	(0.182)	(0.145)
16	0.1	-0.393^{*}	-0.227
	(0.020)	(0.171)	(0.147)
17	0.9	-1.444^{***}	-0.922^{***}
	(0.021)	(0.237)	(0.176)
18	0.4	0.001	0.104
	(0.013)	(0.218)	(0.159)
19	0.1	-0.341^{*}	0.014
	(0.029)	(0.172)	(0.122)
20	0.3	-0.744^{***}	-0.494^{**}
	(0.017)	(0.208)	(0.165)
21	0.7	-1.091^{***}	-1.023^{***}
	(0.012)	(0.253)	(0.215)
22	0.3	-0.751^{**}	$-0.347^{\dagger}$
	(0.025)	(0.276)	(0.195)

Table OS.10 Simulation Study with the Correct State Space $\Pi(s_t)$ and State-Independent State Space $\Pi(\bar{s})$

	Hospital 2			Hospital 13			Hospital 22		
	True value	Correct action set	State-independent action set $\Pi(\bar{s})$	True value	Correct action set	State-independent action set $\Pi(\bar{s})$	True value	Correct action set	State-independent action set $\Pi(\bar{s})$
$\hat{\beta}$	0.300	0.380	0.900	0.500	0.566	0.900	0.900	0.834	0.900
		(0.007)	(0.000)		(0.008)	(0.000)		(0.010)	(0.000)
$\hat{u}_{w,l}$	-0.015	-0.101	0.000	0.000	-0.124	0.000	-1.669	-1.570	0.000
		(0.012)	(0.000)		(0.014)	(0.000)		(0.022)	(0.000)
$\hat{u}_{r,l}$	1.489	1.486	0.248	1.745	1.742	0.364	2.065	2.074	0.010
		(0.011)	(0.007)		(0.001)	(0.007)		(0.003)	(0.005)
$\hat{u}_{w,h}$	-0.750	-0.798	0.000	-0.789	-0.873	0.000	-1.530	-1.475	0.000
		(0.008)	(0.000)		(0.010)	(0.000)		(0.012)	(0.000)
$\hat{u}_{r,h}$	0.301	0.299	-0.863	0.707	0.707	-0.533	0.660	0.665	-1.177
		(0.002)	(0.002)		(0.002)	(0.003)		(0.004)	(0.002)

Table OS.11 Estimated Discount Factor $\hat{\beta}$ Under Alternative Normalization Choices of $(u_{a.l}, u_{a,h})$

Hosp	(0.3,0.3)	(0.2,0.2)	(0.2,0.3)	(0.1,0.2)	(0.1,0.1)	(0,0)	(-0.1,-0.1)	(-0.2,-0.1)	(-0.2,-0.2)	(-0.3,-0.2)	(-0.3,-0.3)
1	0.2	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3
2	0.4	0.5	0.5	0.5	0.5	0.5	0.6	0.6	0.6	0.6	0.6
3	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
4	0.3	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4
5	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
6	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2
7	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9
8	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
9	0.5	0.5	0.5	0.5	0.5	0.5	0.6	0.6	0.6	0.6	0.6
10	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
11	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9
12	0.5	0.5	0.5	0.6	0.6	0.6	0.6	0.6	0.6	0.7	0.7
13	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
14	0.2	0.2	0.2	0.2	0.2	0.3	0.3	0.3	0.3	0.3	0.3
15	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
16	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
17	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9
18	0.3	0.3	0.3	0.3	0.3	0.4	0.4	0.4	0.4	0.4	0.4
19	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
20	0.2	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3
21	0.6	0.6	0.6	0.6	0.6	0.7	0.7	0.7	0.7	0.8	0.8
22	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.4	0.4

Table OS.12 Estimation Results of Dynamic Discrete Choice Model with Time Trend

Hosp	$\hat{\beta}$	$\hat{u}_{w,l}$	$\hat{u}_{r,l}$	$\hat{u}_{w,h}$	$\hat{u}_{r,h}$	$\hat{\beta}$ (main)
1	0.2	0.000	1.445	-0.692	0.304	0.3
2	0.5	-0.023	1.486	-0.682	0.522	0.5
3	0.1	-0.169	2.568	-1.562	1.188	0.1
4	0.4	-0.099	2.322	-1.024	0.839	0.4
5	0.1	-0.223	1.834	-1.209	0.290	0.1
6	0.2	-0.088	1.907	-1.011	0.528	0.2
7	0.6	-0.101	2.731	-0.752	1.605	0.9
8	0.1	-0.277	1.465	-0.794	0.387	0.1
9	0.5	-0.001	1.808	-0.843	0.716	0.5
10	0.1	-0.620	2.350	-1.726	0.932	0.1
11	0.4	-0.056	1.846	-0.908	0.129	0.9
12	0.6	-0.134	2.139	-0.559	0.640	0.6
13	0.1	-0.074	2.235	-1.292	0.810	0.1
14	0.2	0.000	1.930	-0.853	0.673	0.3
15	0.1	-1.105	2.160	-2.280	0.824	0.1
16	0.1	-0.194	2.334	-1.221	1.012	0.1
17	0.9	-1.701	2.070	-1.540	0.660	0.9
18	0.3	-0.001	2.336	-1.244	1.008	0.4
19	0.1	-0.911	1.820	-1.888	0.513	0.1
20	0.2	0.000	2.209	-1.238	0.906	0.3
21	0.6	0.000	1.980	-0.489	0.247	0.7
22	0.3	-0.033	1.546	-1.037	0.119	0.3

Table OS.13 Estimation Results for Individual Hospitals with Three Severity Classes

Hosp	Dim. of state	$\hat{\beta}_{3c}$	$\hat{u}_{w,1}$	$\hat{u}_{w,2}$	$\hat{u}_{w,3}$	$\hat{u}_{r,1}$	$\hat{u}_{r,2}$	$\hat{u}_{r,3}$	$\hat{\beta}_{2c}$ (two class)
2	11616	0.3	0.000	-0.344	-0.750	1.699	1.459	0.301	0.3
3	18144	0.5	0.000	-0.251	-0.659	1.742	1.401	0.518	0.5
5	4680	0.2	-0.229	-0.654	-1.709	3.045	2.250	1.186	0.1
7	22528	0.5	-0.180	-0.733	-1.064	2.870	2.209	0.798	0.4
8	1920	0.2	-0.161	-0.805	-1.244	2.216	1.301	0.333	0.1
9	8316	0.3	-0.130	-0.636	-1.076	2.342	1.715	0.530	0.2
10	2400	0.9	-1.794	-2.032	-1.875	2.881	2.689	1.454	0.9
11	7140	0.1	-0.191	-0.383	-0.763	1.751	1.278	0.380	0.1
13	13248	0.6	-0.068	-0.570	-0.912	2.243	1.658	0.706	0.5
14	3640	0.9	-2.563	-3.034	-2.667	2.740	1.898	0.920	0.1
15	1400	0.9	-1.134	-1.275	-1.366	2.144	1.272	0.105	0.9
16	13475	0.6	-0.061	-0.388	-0.558	2.595	1.897	0.637	0.6
17	4080	0.2	-0.131	-0.545	-1.379	2.641	2.052	0.785	0.1
19	19943	0.2	0.000	0.000	-0.840	2.181	2.006	0.662	0.3
20	3400	0.1	-0.784	-1.850	-2.276	2.531	1.881	0.761	0.1
21	3360	0.1	0.000	-0.621	-1.233	2.688	2.041	1.005	0.1
22	1750	0.9	-1.812	-1.649	-1.530	2.427	1.732	0.660	0.9
24	10395	0.4	-0.021	-0.597	-1.392	2.710	2.106	1.006	0.4
25	8190	0.1	-0.615	-1.447	-1.885	2.202	1.574	0.516	0.1
27	2400	0.3	-0.010	-0.372	-1.333	2.569	1.952	0.861	0.3
29	3780	0.7	-0.104	-0.305	-0.612	2.379	1.594	0.241	0.7
30	3060	0.4	-0.118	-0.309	-1.073	1.865	1.205	0.174	0.3

Table OS.14 **Estimation Results Under Changes in Capacity Choices (All Hospitals Combined)**

	$\hat{\beta}$	$\hat{u}_{w,l}$	$\hat{u}_{r,l}$	$\hat{u}_{w,h}$	$\hat{u}_{r,h}$
Main Specification	0.3	−0.071	1.950	−0.932	0.671
Increase ICU Capacity by One Bed	0.3	−0.066	1.956	−0.928	0.677
Decrease ICU Capacity by One Bed	0.3	−0.090	1.931	−0.958	0.651
Remove Buffer Term In ED Capacities	0.3	−0.070	1.949	−0.931	0.672

Table OS.15 **Estimated Discount Factor $\hat{\beta}$ Under Changes in Capacity Choices**

Hosp	Main Specification	Increase ICU Capacity by One Bed	Decrease ICU Capacity by One Bed	Remove Buffer Term in ED Capacity
1	0.3	0.3	0.3	0.3
2	0.5	0.5	0.5	0.5
3	0.1	0.1	0.1	0.1
4	0.4	0.4	0.4	0.4
5	0.1	0.1	0.1	0.1
6	0.2	0.2	0.2	0.2
7	0.9	0.9	0.3	0.9
8	0.1	0.1	0.1	0.1
9	0.5	0.5	0.5	0.5
10	0.1	0.1	0.1	0.1
11	0.9	0.4	0.3	0.9
12	0.6	0.6	0.6	0.6
13	0.1	0.1	0.1	0.1
14	0.3	0.3	0.2	0.3
15	0.1	0.1	0.1	0.1
16	0.1	0.1	0.1	0.1
17	0.9	0.2	0.1	0.9
18	0.4	0.4	0.4	0.4
19	0.1	0.1	0.1	0.1
20	0.3	0.3	0.3	0.3
21	0.7	0.7	0.7	0.7
22	0.3	0.3	0.3	0.3

Table OS.16 Summary of Model Assumptions and Justifications/Robustness Checks

Assumption	Summary	Justification/Robustness checks
Linear utility function	Linearity is critical to achieve joint identification of discount factor and utility parameters (Komarova et al. 2018). It cannot be directly tested from the data.	It is a reasonable assumption as per period utility should be the sum of action utilities for each patient. In addition, we find in the data that the number of patients in ED is small for most periods (Table OS.18) and the ED waiting time is short for most patients (Table OS.17). We also find that the ICU departure rates are generally similar for different ICU occupancy levels (Table OS.21).
Constant utility parameters	We assume the utility parameters are constant and independent of the state. This is essential crucial for our identification and interpretation of the discount factor in our model. This assumption can be validated empirically given sufficient variation in the data.	The ICU typically allocates dedicated resource to each bed and maintains a fixed nurse-to-bed ratio, thus the benefit of ICU care tends to remain stable even when the ICU is congested. Consistent with this, our reduced-form analysis finds no significant impact of ICU occupancy on patient outcomes (see Table OS.2 in Section OS.3).
Normalization choices	Normalization assumption is needed to identify dynamic discrete choice model (e.g., Magnac and Thesmar 2002). Different normalization choices can affect our estimates of discount factor and utility parameters. The normalization choice cannot be pinned down from the data.	We set the ICU admission utility to zero as (i) it provides a proper reference point for current period utility across ED states; (ii) the ICU has very high fixed and operating costs. We perform robustness checks with stratified sample and alternative normalization choices (see Section EC.7).
Lower ED waiting utility	This assumption is critical in our problem as it rules out unrealistic model estimates (long-term benefit from ED waiting). It cannot be directly tested from the data.	ED waiting represents an additional cost in the patient rerouting process, and is a only temporary state instead of final destination. Thus, on average, it brings lower perceived utility than ICU and non-ICU admission. We find in the data that the ED waiting time is indeed short for most patients (see Table OS.17).
Patient classes	This assumption facilitates our estimation by reducing the state space dimension, which is important as our state space grows exponentially in number of patient classes. It is innocuous as two patient classes are sufficient to capture the difference in patients' needs for care.	We perform robustness checks with three patient classes in Section OS.5. The estimated discount factors and heterogeneity across hospitals are similar to our main specification.
Constant arrival/departure rates	These set-ups are standard in dynamic discrete choice models. The arrival/departure rates are estimated from data.	We perform robustness checks using time-varying transition rates (Section OS.4) and stratified sample analysis (Section EC.7).
ED and ICU capacity	These set-ups are standard in dynamic discrete choice models. The units capacity levels are directly estimated from the data.	The ICU/ED capacity levels are directly estimated from data. We perform robustness checks by varying capacity levels of ED and ICU (Section EC.7).
Unlimited capacity for non-ICU units	This assumption facilitates our estimation by reducing the state space dimension. It can be validated from the observed data.	We observe in the data that the occupancy of non-ICU units seldom reaches their maximum observed level (see footnote 2). Thus, the capacity of non-ICU units should not be a constraint. This is also confirmed in our communication with hospitals.

Table OS.17 Statistics of Patients ED Length-of-Stay

	Full cohort (N = 180051)			ICU admission cohort (N = 20977)		
ED LOS	$\leq$ 2 hours	$\leq$ 4 hours	$\leq$ 6 hours	$\leq$ 2 hours	$\leq$ 4 hours	$\leq$ 6 hours
Proportion	85.40%	94.10%	96.60%	82.40%	93.80%	97.00%

Table OS.18 Summary Statistics of ED Occupancy

	25th percentile	Median	Mean	75th percentile	90th percentile
Low severity class	0.00	1.00	1.02	2.00	3.00
High severity class	0.00	0.00	0.19	0.00	1.00

Table OS.19 Summary Statistics of ICU Length-of-Stay (Hours)

Severity Class	Count	Mean	Std
Low	11490	45.09	48.33
High	8189	66.04	59.98

Table OS.20 Variance-to-mean ratio for arrivals in each periods

Hosp	ED low severity	ED high severity	External ICU
1	1.066	0.985	1.047
2	1.259	1.025	1.072
3	1.025	1.009	1.050
4	1.134	0.999	1.244
5	1.090	1.011	1.038
6	1.085	1.029	1.056
7	1.054	1.012	1.049
8	1.091	0.987	1.043
9	1.083	1.010	1.039
10	1.110	0.994	1.015
11	1.002	1.010	1.016
12	1.316	1.051	1.048
13	1.054	1.009	1.026
14	1.095	0.965	1.175
15	1.077	0.994	1.079
16	1.100	1.012	1.022
17	0.970	0.975	1.020
18	1.171	1.069	1.391
19	1.019	1.003	1.059
20	1.007	0.990	0.996
21	1.003	0.997	1.076
22	1.037	0.976	1.033

Table OS.21 ICU Departure Rate by ICU Occupancy Level (All Hospitals Combined)

ICU Occupancy Decile Bin	ICU Departure Rate μ_I
1	0.031
2	0.031
3	0.032
4	0.032
5	0.031
6	0.029
7	0.032
8	0.032
9	0.030
10	0.029

References

Komarova, Tatiana, Fabio Sanches, Daniel Silva Junior, Sorawoot Srisuma. 2018. Joint analysis of the discount factor and payoff parameters in dynamic discrete choice models. *Quantitative Economics* **9**(3) 1153–1194.

Magnac, Thierry, David Thesmar. 2002. Identifying dynamic discrete decision processes. *Econometrica* **70**(2) 801–816.

Rust, John. 1987. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society* 999–1033.