

Refining Workload Measure in Hospital Units: From Census to Acuity-Adjusted Census in Intensive Care Units

Song-Hee Kim ^{*}, Edieal Pinker [†], Joan Rimar [‡], and Elizabeth Bradley [§]

Abstract

This study aims to better understand the impact of Intensive Care Unit (ICU) workload on patient outcomes, so that practitioners and researchers can use such understanding to provide high quality care despite increased hospital crowding. We use data collected from the medical ICU and the surgical ICU of a major teaching hospital. We measure ICU workload in a novel way that takes into account not only the census but also patient acuity. Having categorized the ICU workload at time of patient departure as low census, high-census/low-acuity, and high-census/high-acuity, we find that patients discharged on a day of high-census/high-acuity ICU workload had significantly worse health status than patients discharged from a low-census ICU workload. Moreover, we find patients with poorer health status at the time of ICU discharge experienced worse longer-term outcomes, including longer post-ICU length-of-stay (LOS), higher mortality, and higher total hospital costs. In sum, we find that acuity was critical in accurately characterizing ICU workload, resulting ICU discharge decisions, and ultimately patient outcomes and hospital costs. Our findings suggest that (1) ICUs need to track the changes in patient acuity in addition to census, in order to use such information to prevent reaching high census with high acuity patients and (2) future studies (both empirical and analytical) of ICU workload should take patient acuity into account in ICU workload measures. Lastly, using a simulation study, we show how high-census/high-acuity workload can be prevented by reducing seasonality in the volume and acuity of patient arrivals.

Keywords: healthcare delivery, intensive care units, acuity-adjusted workload, congestion, Rothman Index, acuity, empirical operations management

Date: August 15, 2016

1 Introduction

Traditionally, patient census (i.e., the number of patients in beds) has been the sole tool to determine the workload in hospital units. For staffing and sizing hospital units, the midnight census is the standard indicator of workload (Beswick et al. 2010). Studies have examined the impact of workload on patient care and outcomes, and most of them have measured workload by census only (Iwashyna et al. 2009, Kc and Terwiesch 2012, Kim et al. 2015). Performance analysis models for hospital units developed by the operations

^{*}Marshall School of Business, University of Southern California, Email: songheek@marshall.usc.edu

[†]Yale School of Management, Yale University, Email: edieal.pinker@yale.edu

[‡]Strategic Analytics, Yale New Haven Health System, Email: joan.rimar@ynhh.org

[§]Yale School of Public Health, Yale University, Email: elizabeth.bradley@yale.edu

research/management community also have measured workload using census only (Green 2002, Dobson et al. 2010, Chan et al. 2012, Shi et al. 2015). However, the acuity of patients in a hospital often varies widely, especially in Intensive Care Units (ICUs), which provides the highest level of care for critically ill patients. Higher acuity patients have been shown to require more time and attention of the physicians and nurses and to generate more workload (Welton 2007, Miranda and Jegers 2012).

It is said that “What gets measured is what is improved”. This paper aims to explore whether adjusting for patient-acuity is important in accurately characterizing ICU workload. Specifically, we categorize the ICU workload at time of patient departure as low census, high-census/low-acuity, and high-census/high-acuity, and study its impact on patient outcomes. ICUs usually have high utilization, stochastic arrivals, and stochastic patient healing processes, which makes it inevitable at times for the demand for ICU care to exceed the capacity of the ICU (Iwashyna et al. 2009). Hence, advancing the understanding of the impact of heavy ICU workload on patient outcomes will help practitioners provide high quality care while managing high demand for ICU care. Researchers will also be able to develop more realistic and effective models.

Patient outcomes we consider include the health status of patients at the time of discharge from the ICU. Most of the existing studies have not examined the health status of patients, at the point of ICU discharge, as a patient outcome, mostly because of data limitations. They have instead examined longer-term outcomes, such as subsequent hospital mortality (Tarnow-Mordi et al. 2000, Iwashyna et al. 2009), ICU or hospital readmission rates (Chrusch et al. 2009, Kc and Terwiesch 2012), post-ICU hospital length-of-stay (LOS) (Wagner et al. 2013), and hospital discharge destination (Wagner et al. 2013). While longer-term outcomes are important, using longer-term outcomes as a primary measure of the impact of ICU workload on patient outcomes gives an incomplete picture because they are inevitably influenced by the operations of the hospital and the downstream units. In addition, some studies have examined proxies for patient health status, such as ICU LOS (Kc and Terwiesch 2012). However, the link between LOS and health status is complex, and longer LOS does not necessarily result in improved health status, as we illustrate with an example in Section 2.

We use patient data collected from two ICUs—a Medical ICU and a Surgical ICU—of a large urban U.S. teaching hospital. To examine patient health status, we take advantage of a relatively new measure of patient acuity called the *Rothman Index* (RI) (Rothman et al. 2013). The novelty of the RI is that in hospitals that implemented the RI system, RIs are automatically calculated from the electronic medical record data and are updated frequently throughout a patient hospital stay.¹ In our data, patient RIs were updated every hour while in ICUs, allowing us to track the health status of patients very close to the time of their discharge from an ICU

¹ A recent Wall Street Journal article discusses the use of the RI: Landro, L. (2015, May 25). Hospitals Find New Ways to Monitor Patients 24/7. Wall Street Journal. Retrieved from <http://www.wsj.com>.

and to examine its direct association with workload. RIs also enable us to measure the acuity-adjusted ICU workload, because at any point of time in our study period, we know not only the census but also the health status of each of the patients.

We find that patients discharged on a day of high-census/high-acuity ICU workload had significantly worse health status than patients discharged from a low-census workload ICU. To our knowledge, this is the first study to show that more acutely ill patients are discharged from an ICU when the acuity-adjusted workload is high rather than low. Moreover, we find patients with poorer health status at the time of ICU discharge experienced worse longer-term outcomes, including longer post-ICU length-of-stay (LOS), higher mortality, and higher total hospital costs.

Our findings provide compelling evidence that acuity is critical in accurately characterizing ICU workload, resulting ICU discharge decisions, and ultimately patient outcomes and hospital costs. The findings suggest that ICUs need to track the changes in patient acuity in addition to census, in order to use such information to prevent reaching high-census/high-acuity state. Using a simulation study, we show how high-census/high-acuity workload can be partially prevented by reducing seasonality in the volume and acuity of patient arrivals. While reducing seasonality in patient arrivals may be impossible in Medical ICUs, where patient arrivals are not scheduled, we find clear day-of-week seasonality in patient arrivals in the Surgical ICU we study, among the patients coming from the operating room after elective surgeries. We show that scheduling the surgeries more uniformly (in terms of both the number and the acuity of the patients) across the days of the week could lead to being in high-census/high-acuity state less often. When high-census/high-acuity is unavoidable, hospitals could improve the quality of care in downstream units so that a patient prematurely discharged from an ICU is more likely to recover without a need to return to the ICU. Our findings also suggest that future studies (both empirical and analytical) of ICU workload should take patient acuity into account in ICU workload measures.

The rest of this paper is organized as follows. In Section 2 we explain our main contributions relative to the relevant existing literature. In Section 3, we describe our dataset, the two ICUs we examine, and the sample we use for our analyses and introduce our new measure of workload. In Section 4, we present our hypotheses and econometric model and show the impact of acuity-adjusted workload on patient health status. In Section 5, we show the impact of patient health status on longer-term patient outcomes including health status upon hospital discharge, post-ICU LOS, ICU readmission rate, 30-day hospital readmission rate, in-hospital mortality rate, and total hospital costs. In Section 6 we discuss possible managerial interventions and use a simulation model to show the potential benefit of smoothing arrivals to an ICU. We conclude and present future research directions in Section 7.

2 Literature Review

In the operations management literature, there has been an increasing interest in empirically studying the impact of workload on system performance. Many of the studies were conducted in hospital settings. For example, studies examined the impact of emergency room workload on nurse absenteeism rates ([Green et al. 2013](#)), on ambulance diversion ([Allon et al. 2013](#)), and on service times ([Kuntz and Sülz 2013](#), [Batt and Terwiesch 2016](#)). By examining a patient transport service system and a cardiothoracic surgery system, [Kc and Terwiesch \(2009\)](#) studied the impact of workload on service times and quality of care. Freeman and colleagues ([2016](#)) examined the impact of workload on resource use and cost of care in a maternity unit. Studies also examined the impact of workload at the hospital level or the hospital department level on the accuracy of hospital discharge coding ([Powell et al. 2012](#)), on in-hospital mortality ([Kuntz et al. 2015](#)), and on service times ([Jaeker and Tucker 2016](#)). Outside of the healthcare setting, Hasija and colleagues ([2010](#)) examined the impact of workload in an email contact center on agent service times and [Tan and Netessine \(2014\)](#) examined the impact of workload in a restaurant chain on server sales efforts.

More closely related studies to this paper examined the impact of workload in ICU settings. Studies in both the operations management literature and the medical literature have shown that ICU workload is an important factor affecting ICU admission decisions. That is, patients, who would have received ICU care otherwise, were not admitted to the ICU because of high workload and limited capacity (e.g., see [Singer et al. \(1983\)](#), [Robert et al. \(2012\)](#) and [Kim et al. \(2015\)](#)). Given that ICU care must be “rationed” at such times, some studies investigate how the rationing could be done in the most effective way. For instance, Shmueli and colleagues ([2003](#)) and Kim and colleagues ([2015](#)) proposed ICU admission policies and evaluated their performance on resulting patient outcomes including patient mortality and subsequent hospital LOS. Dobson and colleagues ([2010](#)) developed a performance analysis model for ICUs with patient blocking and bumping for given arrival patterns and capacity. We note that we do not examine the impact of workload on patient admission to an ICU in this paper; we study the impact of ICU workload on the care provided to patients *after* ICU admission.

Studies that examined the impact of ICU workload on patients already in the ICU, then, are the most relevant studies to this paper. [Kc and Terwiesch \(2012\)](#) showed that patients discharged when an ICU had high census had a shorter ICU LOS, which in turn led to a higher chance of getting readmitted to the ICU later in their hospital stay. Their finding implied that congestion in the ICU shortened the care for some patients, which led to worse patient outcomes in the sense that these patients were more likely to need further ICU care. On a similar note, Anderson and colleagues ([2011](#)) investigated daily discharge rates from a surgical ICU and found higher discharge rates on days with high census and more scheduled surgeries. This “premature”

discharge from an ICU due to high workload has been recognized as a problem in many ICUs, and Chan and colleagues (2012) has for example developed a model to attempt to optimize the choice of which patients to discharge to reduce health risks. On the other hand, Long and Mathews (2015) separated the ICU LOS into service time and time-to-transfer, and showed that high census led to a shorter time-to-transfer, but did not affect the service time. Their findings suggest that workload does not influence the care provided to the patients.

The medical literature has more mixed evidence for the impact of workload in ICUs. For example, Tarnow-Mordi and colleagues (2000) found that ICU patients exposed to high workload were more likely to die than patients exposed to low ICU workload. Chrusch and colleagues (2009) found that occupancy based markers of unit activity are associated with the likelihood of early death or ICU readmission. In contrast, Iwashyna and colleagues (2009) found that patients admitted on high census days had the same odds of in-hospital mortality as patients admitted on average or on low census days. Also, Wagner and colleagues (2013) found no association between ICU workload and long-term outcomes, including subsequent in-hospital mortality, post-ICU discharge LOS, and hospital discharge destination. The conflicting results could be because of data limitations in the previous studies, which did not allow adjusting for patient acuity in workload measures. In addition, previous studies have included patients who were so sick that they were unlikely to benefit from ICU care and have much higher mortality rates than other patients. Because decisions about when to discharge these patients from the ICU are made according to different criteria than patients who are benefiting from ICU care, mixing all these patients together can distort the results. In this study, we remove the “too-sick-to-benefit” patients. See footnote 3 for details.

In relation to the literature mentioned above, we make three main contributions:

First, we introduce a new workload measure that takes into account patient acuity in addition to census, and we show that it is a better workload measure than census alone. We show, by classifying high census days using the percentage of acute patients in the ICU, that the effects of high census on the health status of patients discharged from an ICU are driven by the cases with a higher percentage of acute patients. Previous literature suggests that hospitals might want to avoid reaching high occupancy levels to prevent their adverse influence on patient care and outcomes. Our results suggest that it is high occupancy levels *with* many high acuity patients that hospitals want to avoid, because having high occupancy levels with many low acuity patients might not be as problematic. Using a simulation model we show that this effect can be explained by sampling. That is, when an ICU is crowded and space is needed, patients who are healthier relative to other patients will be discharged. If the ICU has many high acuity patients the discharged patients will tend to be sicker even if they are healthier relative to the other patients in the unit.


We note that Wagner and colleagues (2013) also included patient acuity in their ICU workload definition; they averaged the acuity of patients in the ICU, based on individual acuity of illness scores calculated on the day of ICU admission. Because patient condition changes throughout ICU stays (e.g., see Figure 1), we believe our workload measure based on the hourly-updated condition of patients in the unit gives a more accurate measure of ICU workload. Furthermore, Wagner and colleagues (2013) used an average (across patients in the unit) acuity score and a separate census variable in their model. Such specification is unable to separate the effect of high census with low acuity from that of high census with high acuity. In addition, Tarnow-Mordi and colleagues (2000) used patient census and nursing requirement² suggested by the UK Intensive Care Society based on patient condition to define workload. However, nursing workload is only a partial view of the workload in an ICU; for instance, nursing workload does not take into account physician activities and it is derived using rough categories of care and nurse to patient ratios. We believe our new measure of ICU workload is a more accurate measure of ICU workload compared to the existing measures.

Second, we show the *direct* relationship between high workload in an ICU and the health status of patients being discharged from such an ICU. Specifically, we find that more acutely ill patients are discharged from the ICU when the workload is high rather than low. Previous studies examined the relationships between high workload in an ICU and *proxies* for health status or *longer-term patient outcomes*. For instance, Kc and Terwiesch (2012) considered *ICU LOS* as a proxy for health status. However, the link between LOS and health status is complex. As an example, Figure 1 shows the changes in RI of two medical ICU patients in our data. The patient on the left came into the ICU with a low RI (lower RI indicates poorer condition—see Section 3.1 for a detailed description of RI), stayed for 4 days, and left with a high RI, whereas the patient on the right came into the ICU with a higher RI, stayed for 11 days, and left with a low RI. In other words, the sicker patient (the patient on the left) stayed a shorter time in the ICU but left healthier than the other patient (the patient on the right) who came in healthier. This example illustrates that a longer stay does not necessarily result in improved health status, and thus highlights the need to look beyond the impact of workload on LOS and to examine the direct association between workload and health status.

Chrusch and colleagues (2009) and Kc and Terwiesch (2012) considered *ICU readmission (bounce back) rates*. While reducing ICU readmission rates is a worthy goal, using these rates as a primary measure of the impact of ICU workload on patient flow gives an incomplete picture for the following reasons: (i) there may be patients who are not readmitted to the ICU because of heroic efforts in downstream units. The ability of downstream units to prevent ICU readmissions may also vary from one downstream unit to another and

²In fact, there have been efforts to measure nursing workload in the intensive care unit, including the Therapeutic Intervention Scoring System (TISS) and the Nursing Activities Score (NAS); see Miranda et al. (2003) for details. An instrument for measuring physician activities is not yet available (Miranda and Jegers 2012).

Figure 1: Rothman Index Scores of Two Medical ICU Patients



from one hospital to another, making it difficult to generalize observations about the effect of workload on ICU patient readmissions; (ii) there may be ICU readmissions that we do not observe because the ICU is crowded and can not accept them back; (iii) patients who are discharged sicker because of workload in the ICU will generally put a strain on downstream units, and only looking at ICU readmission rates hides this; and (iv) some people die on general units unexpectedly after ICU discharge (i.e., staff did not recognize the impending death), and they may not have been considered for ICU readmission.

Wagner and colleagues (2013) considered *in-hospital mortality rates*. Similar to ICU readmission rates, in-hospital mortality rates are also influenced by the operations of the downstream units. Furthermore, (fortunately) only a small portion of the patients die in hospital, and ICU workload is likely to affect even a smaller portion of them. In contrast, we find a negative health impact on patients even when excluding the patients who subsequently died in the hospital. That is, this is a measurable impact that affects a broad range of patients.

Previous studies could not examine the direct relationship of workload and health status because they used static acuity measures that fail to track the changes in patient condition. For example, Kc and Terwiesch (2012) used the Euroscore (2007), which measures the risk of death based on patient information before a heart operation and hence before ICU admission. The Euroscore was used as a proxy for patient acuity at ICU discharge, but Figure 1 illustrates that patients who are the sickest of new arrivals are not necessarily the sickest among those being discharged. (To illustrate this point systematically, we show RI quartile rankings for patients at admission and discharge for the Medical ICU and the Surgical ICU in Table 3.) Indeed, most of the acuity scoring systems within ICUs—including the Acute Physiology and Chronic Health Evaluation (APACHE), Simplified Acute Physiology Score (SAPS), and Mortality Prediction Model (MPM)—are static, and are often computed only once, within 24 hours of admission of a patient to an ICU (Strand and Flaatten 2008).

In contrast, the RI tracks patient condition changes, and allows us to track the acuity of the patients at discharge from the ICU. Chrusch and colleagues (2009), in fact, acknowledge such a need to look at health status at discharge (on page 2756): “Consideration could be given to using a score to further adjust for patient acuity at the time of discharge, such as acute physiology scores, nursing workload measures, or organ dysfunction scores.” To our knowledge, we are the first to examine the *direct* relationship between high workload in an ICU and the health status of patients being discharged from such an ICU.

Third, we discuss a possible intervention for an ICU to moderate the impact of high acuity-adjusted workload on patient care within the context of other possible managerial actions. Generally the medical literature on workload in ICUs does not discuss potential managerial interventions. In the operations management literature, Kc and Terwiesch (2012) claimed that ICU readmissions end up using more capacity than was saved by premature discharges. They concluded with the recommendation that ICU managers not discharge aggressively and instead “do it right the first time”. This prescription implies that ICU managers actually have a choice. We disagree because we do not see evidence that the ICU is being “aggressive” in discharging and do not know how aggressiveness would be identified. Our analysis indicates that ICU physicians are tending to discharge the healthiest patient first and only under high occupancy with many high acuity patients do they discharge patients who might be more acute than if under low occupancy. Using a simulation we explore how seasonality in patient volume and acuity both can cause the ICU to operate more frequently under high occupancy with many high acuity patients, leading to more early discharges. We identify demand smoothing as a more realistic mechanism to reduce the negative impact of high workloads.

3 Setting

3.1 Data

We collected hospitalization data for every adult patient who received care in the Medical ICU (MICU) or the Surgical ICU (SICU) at a major U.S. teaching hospital over the course of 15 months. For each hospitalization, we had the patient age, gender, principal and secondary diagnoses (identified by the ICD-9 codes), principal and secondary procedures (identified by the ICD-9 codes) and their dates, Diagnosis-related Group (DRG) classification, payor, in-hospital death, and discharge dispositions. We generated 30 comorbidity indexes (e.g., diabetes) using the ICD-9 codes and the DRG classification (Elixhauser et al. 1998). The data also included every unit the patient visited along with unit admission and discharge dates and times, including hospital admission and discharge date and time.

In addition, a key feature of our data was the availability of the patient *Rothman Index* (RI) scores throughout each hospitalization. The RI score is a composite measure updated regularly from the electronic medical

record based on changes in 26 clinical measures including vital signs, nursing assessments, Braden score, cardiac rhythms, and laboratory test results; see [Rothman et al. \(2013\)](#) for details. This score is independent of diagnosis, and it was developed to be used for any inpatient (i.e., medical or surgical patients including critical care patients). With a theoretical range from -91 to 100, the majority of patients on a general medical or surgical unit fall within the range from 0 to 100, with lower scores indicating poorer condition. Studies have shown that the RI score is associated with 24-hour mortality, 1-year mortality, APACHE III score, and discharge disposition ([Rothman et al. 2013](#)); 30-day hospital readmission rates ([Bradley et al. 2013](#)); post-operative complications for colorectal surgery ([Tepas et al. 2013](#)); unplanned ICU transfers ([Danesh et al. 2012](#)); and unplanned SICU readmissions ([Piper et al. 2014](#)). The RI score has also been shown to outperform the Modified Early Warning Score (MEWS)—widely used in hospitals to early detect clinical deterioration—in identifying patients who are likely to die within 24 hours ([Finlay et al. 2014](#)).

We note that during our data collection period, physicians could access patients’ RI scores by clicking a few buttons on the electronic medical record system of the hospital. Because the RI scores were recently introduced to this hospital, the physicians did not regularly view the RI during our data collection period, and their admitting, observation, and discharge decisions were not influenced by the RI scores.

3.2 The Two Intensive Care Units

The MICU had 36 staffed beds, and the SICU had 21 staffed beds during our study period. The initial data include 3,728 MICU visits and 1,845 SICU visits during the initial 15 months period. We utilize patient flow data from all of these visits to compute the census changes in the MICU and the SICU, respectively. We then restrict our study to the one year in the center of the period, from 2/1/2013 to 1/31/2014, to avoid censored estimation of census.

Table 1 provides the summary statistics for daily arrivals, discharges, and workload for the two ICUs in our one year study period. The MICU has an average of nearly 9 new arrivals (and discharges) each day and the SICU has approximately 4. The largest source of MICU patients is the emergency room (about 50%) whereas for the SICU it is the operating room (about 40%). After ICU care, most of the patients—85% of the MICU patients and 88% of the SICU patients—are transferred out to step-down units, general wards, or discharged from the hospital.

Because we know every patient’s visited units and RI trajectories, we can compute not only the census of the MICU and the SICU at any point in time, but also the RIs of the patients in the same unit at any point in time. We measure the daily *workload* of the MICU and of the SICU using the following two variables: (1) the daily average *census* and (2) the daily average percentage of *more acutely ill* patients. We define a patient

Table 1: Summary Statistics for Daily Arrivals, Discharges, and Workload in the Medical ICU (MICU) and the Surgical ICU (SICU). N = 365 days.

	MICU				SICU			
	Mean (Std)	Median	75th pct	Max	Mean (Std)	Median	75th pct	Max
# Daily Arrivals	8.9 (2.8)	9	11	18	4.3 (2.2)	4	6	13
from operating room	0.1 (0.3)	0	0	2	1.7 (1.4)	1	2	8
from emergency room	4.4 (2.1)	4	6	12	1.3 (1.1)	1	2	6
from step-down/wards/other ICUs	3.8 (1.7)	4	5	9	1.2 (1.1)	1	2	6
direct	0.7 (0.9)	1	1	5	0.2 (0.4)	0	0	2
# Daily Discharges	8.9 (2.7)	9	11	16	4.3 (2.0)	4	6	12
to step-down/wards/hospital discharge	7.6 (2.4)	7	9	14	3.8 (1.9)	4	5	11
in-unit death	1.0 (1.0)	1	2	5	0.2 (0.5)	0	0	3
to another ICU	0.3 (0.6)	0	1	4	0.2 (0.5)	0	0	3
to operating room*	0.04 (0.20)	0	0	2	0.04 (0.21)	0	0	2
Daily Avg Census	32.1 (3.3)	32.8	34.6	39.5	16.3 (2.7)	16.9	18.3	20.8
Daily Avg % Patients with $RI \leq 50$	80.9 (6.2)	81.0	85.5	99.1	76.2 (9.2)	76.8	82.9	97.4

* This number includes patients who went to the operating room but was later readmitted to the ICU post-operatively. We note that clinicians would not consider such cases as ICU discharges.

to be ‘more acute’ if her RI is less than or equal to 50, which is approximately the average RI upon ICU admission of patients who are visiting the MICU or the SICU for the first time during their hospitalizations (see Table 2).

Table 1 shows that on average, about 32 of the 36 staffed beds in the MICU are occupied and about 17 of the 21 staffed beds in the SICU are occupied. Furthermore, for 25% of the days in our study period, the MICU has about 34 or more patients and the SICU has about 18 or more patients. We have learned that extra licensed beds in the two ICUs allow the ICUs to temporarily keep more patients than the number of the staffed beds; the maximum values of the ‘Daily Avg Census’ suggest that this does happen, especially in the MICU.

We compute ‘Daily Avg % Patients with $RI \leq 50$ ’ by tracking changes in the census and the number of patients with $RI \leq 50$. For instance, if there are 32 patients in the MICU at 8 am and 24 of them have their $RI \leq 50$, the percentage of more acute patients at 8 am is 75%. If one patient’s RI increases from 45 to 55 at 8:40 am with other things being equal, the percentage of more acute patients drops to 72% at 8:40 am. We then compute the time-weighted average of these percentages over each day to get the ‘Daily Avg % Patients with $RI \leq 50$ ’. From Table 1, we observe that its average is 81% in the MICU and 76% in the SICU. That is, on

average, about 4 out of 5 patients are ‘more acute’ than an average first-time ICU patient. The proportion of ‘more acute’ patients is high because (1) patients who are more acute tend to stay longer in the ICU, affecting the time-weighted average and that (2) patients who are repeat ICU patients (i.e., patients who were in the ICU earlier in the hospitalization and were readmitted to the ICU) tend to be more acute than an average ICU patient who is in the ICU for the first time.

3.3 Sample Selection and Summary Statistics

In examining the impact of workload on ICU care, we attempt to eliminate possible confounding events by restricting our sample as detailed in Table 10 in the appendix. (However, all patients in our data are considered when we compute the workload measures.) Our one year study period has 3,227 MICU visits and 1,556 SICU visits. A patient can have more than one ICU visit during a hospitalization, and we focus on only the first ICU visit. We then exclude transfers from neighboring hospitals, because we do not have enough information on what care the patient received at the previous hospital or during the transfer. We exclude MICU and SICU visits that are followed by an admission to a different ICU; the hospital we study has three additional ICUs—a cardiac ICU, a cardiac-thoracic ICU, and a neuro ICU. We have learned from the clinicians that these patients usually have multiple complications, and that such patients are transferred to other specialty ICUs after the MICU or the SICU to receive additional ICU care. We exclude visits whose next units are the operating room, interventional radiology, catheterization and electrophysiology (Cath/EP), or hemodialysis, because deciding when to send a patient to these next units are unlikely to be affected by the ICU workload. We exclude visits whose length-of-stays are outliers: shorter than 12 hours or longer than 21 days. We also exclude visits preceded by hospital stays longer than 7 days, because such ICU visits are likely to be due to hospital-acquired complications that might not be well explained by our data. We exclude visits with missing RI or patient information. Lastly, we want to exclude the visits of the “too-sick-to-benefit” patients because the way workload affects their care must be different from the way workload affects the “normal” patients³; we remove patients who die later in the hospitalization or patients who are classified “comfort-measures-only

³Among the 1,976 first-time MICU patients who were not discharged home directly after ICU stay, the average *RIX* was 50.5, readmission rate 9.0%, and in-hospital mortality 5.7%; among the 148 first-time MICU patients who were not discharged home directly after ICU stay *and* were classified as a CMO patient during hospitalization, the average *RIX* was 29.9, ICU readmission rate 30.4%, and in-hospital mortality 56.1%. Similarly, among the 1,010 first-time SICU patients who were not discharged home directly after ICU stay, the average *RIX* was 55.8, ICU readmission rate 7.0%, and in-hospital mortality 1.9%; among the 28 first-time SICU patients who were not discharged home directly after ICU stay *and* were classified as a CMO patient during hospitalization, the average *RIX* was 33.7, readmission rate 32.1%, and in-hospital mortality 57.1%. The noticeable differences in the average *RIX*, ICU readmission rate and in-hospital mortality of the CMO patients suggest that the hospital might have different decision processes for such sick patients.

(CMO)” before hospital discharge.

Table 2: Summary Statistics of Patient Health Status, ICU Care, Workload, and Longer-term Patient Outcome Measures for the Medical ICU (MICU) and the Surgical ICU (SICU) Patients

	MICU					SICU				
	N	Mean (Std)	Min	Median	Max	N	Mean (Std)	Min	Median	Max
RIE_i (RI at ICU Arrival)	1864	45.4 (20.9)	-19.4	44.5	97.1	1036	51.8 (18.5)	-15.2	51.3	97.4
RIX_i (RI at ICU Discharge)	1864	53.0 (20.0)	-10.2	52.8	96.8	1036	57.5 (16.4)	8.1	57.6	98.3
$I(RIX > RIE)_i$ (Whether $RIX_i > RIE_i$)	1864	66.1%				1036	63.3%			
$ICULOS_i$ (ICU LOS in days)	1864	2.9 (2.7)	0.5	1.9	20.2	1036	3.1 (3.0)	0.5	2.0	20.0
$I(SD)_i$ (Whether next unit is a step-down)	1864	8.8%				1036	18.1%			
$CensusX_i$ (Avg census on discharge day)	1864	31.8 (3.1)	18.3	32.4	39.0	1036	16.1 (2.4)	6.0	16.7	20.4
$MoreAcuteX_i$ (Avg % more acute on discharge day)	1864	80.8 (6.3)	61.0	81.0	97.2	1036	75.9 (9.2)	46.5	76.6	100.0
$RIHX_i$ (RI at hospital discharge)	1796	69.7 (18.2)	-0.8	73.4	98.9	978	75.6 (14.4)	5.8	78.9	98.5
$LOS_{afterICU}_i$ (Hospital LOS after ICU in days)	1798	7.2 (11.5)	0.1	4.1	272.9	979	6.5 (7.7)	0.3	4.0	81.8
$I(reICU)_i$ (whether readmitted to ICU)	1798	6.6%				979	6.1%			
$I(reHosp)_i$ (whether readmitted to hospital in 30 days)	1864	27.3%				1036	16.7%			
$I(death)_i$ (whether died in-hospital after ICU)	1976	5.7%				1010	1.9%			
$Cost_i$ (Hospital cost in thousand dollars)	1864	30.0 (34.6)	0.2	19.9	643.8	1036	41.5 (35.0)	0.4	30.8	309.4

Our final sample consists of 1,864 MICU visits and 1,036 SICU visits, whose summary statistics for the variables of interest are provided in Table 2 (see Table 11 in the appendix for the summary statistics of other variables). The mean RI at ICU arrival (RIE_i —“E” representing entry) is 45.4 in the MICU and 51.8 in the SICU, whereas the mean RI at ICU discharge (RIX_i —“X” representing exit) is 53.0 in the MICU and 57.5 in the SICU. Approximately 65% of the patients in both the MICU and the SICU have higher RI upon ICU discharge compared to the RI at ICU arrival; in other words, about 65% of the patients have improved health status after the ICU stay. Patients stay 2.9 days in the MICU and 3.1 days in the SICU on average; the ICU LOS distributions have heavy tails and the medians are approximately 2 days for both ICUs. Approximately 9% of the MICU patients and 18% of the SICU patients are transferred to a step-down (SD) unit upon ICU discharge, instead of being transferred to a general ward or being discharged from the hospital. This difference in the use of step-down units by the MICU and the SICU could be due to the inherent difference in the care for medical versus surgical patients as well as due to the number of available step-down beds; there are 15 step-down beds where the MICU patients are usually sent to and 23 step-down beds for the SICU patients.

Table 2 also provides the summary statistics of our workload measures that patient i experienced. We measure the daily average census on patient i ’s discharge day ($CensusX_i$ —“X” representing exit) and the daily average percentage of more acute patients on patient i ’s discharge day ($MoreAcuteX_i$ —“X” again representing exit). For instance, if a patient is discharged from the ICU on August 2, we track the changes in the census and the percentage of more acute patients during the 24 hours between 8/2 12:00 AM and 8/3 12:00 AM and compute the averages. $CensusX_i$ and $MoreAcuteX_i$ could be interpreted as the workload

patient i is exposed to on her departure day. Note that we do not account for patient i 's occupancy or RI in computing $CensusX_i$ and $MoreAcuteX_i$. Hence, two patients discharged on the same day have different values for $CensusX_i$ and $MoreAcuteX_i$.

Lastly, the summary statistics of longer-term patient outcome measures are reported in Table 2. In Section 6, we will examine the association between RI at ICU discharge (RIX_i) versus the six longer-term patient outcomes: RI at hospital discharge ($RIHX_i$), hospital LOS after the ICU stay ($LOS_{afterICU_i}$), ICU readmission during the same hospitalization ($I(reICU)_i$), hospital readmission within 30 days from hospital discharge ($I(reHosp)_i$), in-hospital death during the same hospitalization ($I(death)_i$), and total hospital cost ($Cost_i$). Because patients who are directly discharged from the ICU have $RIHX_i = RIX_i$, $LOS_{afterICU_i} = 0$, $I(reICU)_i = 0$ and $I(death)_i = 0$, we study these four patient outcome measures over the subset of patients whose subsequent LOS after ICU is not zero (i.e., $LOS_{afterICU_i} > 0$). Note that all of the variables in Table 2 are computed over the patients who did not have in-hospital death, except for $I(death)_i$; when computing the mean $I(death)_i$, we include the “too-sick-to-benefit” patients (patients who die later in the hospitalization and patients who are classified comfort-measures-only before hospital discharge).


3.4 Categorizing ICU Workload

In Section 3.3 we showed how we measure the ICU workload in two different ways using $CensusX_i$ and $MoreAcuteX_i$. We use these two variables to categorize ICU workload used throughout this paper. First, we employ a threshold of 34 patients and 18 patients for the MICU and the SICU, respectively, to divide into ‘low’ census and ‘high’ census (this approach has been frequently used by other studies, including Kc and Terwiesch (2012) and Kim et al. (2015)). We then further divide the high census state by whether the percentage of more acute patients is more than 80% or not. Figure 2 illustrates this idea: patients in I experienced low census on discharge day, patients in II experienced high census but smaller percentage of more acute patients (high-census/low-acuity), and patients in III experienced high census and higher percentage of more acute patients (high-census/high-acuity). Among the 1,864 MICU patients, 74.9% were discharged from low-census, 11.8% from high-census/low-acuity, and 13.3% from high-census/high-acuity. Among the 1,036 SICU patients, 77.9% were discharged from low-census, 12.6% from high-census/low-acuity, and 9.5% from high-census/high-acuity.

Previous studies have found evidence that ICU care differs in low-census versus high-census (e.g., see Kc and Terwiesch (2012)). Throughout this paper, we will make that comparison, but in addition will examine high-census/low-acuity and high-census/high-acuity separately. We have tried various other thresholds to

categorize the workload; while different thresholds tell similar stories, the thresholds we choose give us a consistent definition across our two ICUs and also leaves enough data points in each category for effective estimation.

Figure 2: Scatter plots of $CensusX$ and $MoreAcuteX$ and categorizing ICU workload—I: low-census, II: high-census/low-acuity, and III: high-census/high-acuity—in the Medical ICU (MICU, left) and the Surgical ICU (SICU, right)



4 Impact of Acuity-Adjusted ICU Workload on Patient Health Status

4.1 Hypotheses Development

In order to develop intuition about the impact of workload on patient discharge patterns from an ICU we propose a simple model of the operation of the unit and use discrete event Monte-Carlo simulation to generate hypotheses of what we should expect to see in our data if the model is correct. The main operating premise of our model is that the patient health status is completely determined by the RI, that patients are discharged according to their RI which providers can perfectly discern and that new arrivals are given priority over current occupants of the ICU. None of these are completely true but give us a useful starting point for thinking about how workload affects patient flow through an ICU.

Our model operates as follows. Each day a patient stays in an ICU, the patient's RI follows a Markov process which changes by an independent random draw from a distribution with positive mean (positive values imply improvement and negative values imply a deterioration in patient health status). When a patient's RI exceeds a particular threshold they are ready for discharge to another less intensive unit of the hospital. We will call such patients: healthy. The implication of these assumptions is that a patient's LOS is tightly linked to the patient's RI at arrival to the unit. At the start of each day the care providers discharge all patients

who are healthy. Then a random number of new patient arrivals is determined. If the number of new arrivals exceeds the number of available beds then the providers will discharge additional patients from the ICU who are not yet healthy. These additional (premature) discharges will be done in order of healthiest (highest RI) patients first. Each new arrival has an initial RI that is drawn from a distribution.

If the ICU operates as described we expect to see the following:

Hypothesis 1: Patients who are discharged when workload is low (low-census) will tend to be healthier (have higher RI) than patients discharged when the workload is high (high-census/low-acuity and high-census/high-acuity).

This is because when workload is high you are more likely to need to discharge a patient who is not yet healthy thus they have a lower RI.

Hypothesis 2: Similarly, the LOS of patients who are discharged when workload is low (low-census) will tend to be longer than those discharged when the workload is high (high-census/low-acuity and high-census/high-acuity).

Our model assumes that the care providers will discharge the healthiest patients available from among the current occupants even when making premature discharges. This will mitigate the impact of workload on the reduction in RI and LOS of discharged patients. We expect to see that mitigation by distinguishing between cases when a patient was discharged from high-census/high-acuity high-census/low-acuity. That is, we expect the following:

Hypothesis 3: Patients discharged in high-census/high-acuity will have lower RI and LOS than patients in low-census and high-census/low-acuity.

We conduct simulation experiments using parameter values derived from our data, and find that these hypotheses are consistent with our simulation results; Section A of the appendix provides details of the simulation experiments and their results. In practice, however, the simple operational model of an ICU behind our simulation study has a number of flaws and thus we cannot rely on it for estimates of the impact of workload on patients. First, we observe some evidence from our data that admissions to the MICU decline when workload is high (We note that [Kim et al. \(2015\)](#) had a similar finding). Second, in practice a patient's readiness to be discharged from the ICU is not completely determined by the RI, and other factors (such as diagnosis) are considered in the ICU discharge decision process. Third, some patient's health status plateaus in the ICU representing neither an improvement from a low level or a stabilization of a deterioration from a higher level. In either case they do not gain any additional value from staying in the ICU and are thus discharged even though their health status and RI is lower than others still in the ICU, violating the simulation model assumption that only the healthiest patients are discharged.

4.2 Econometric Model

We now describe the model we use to test our hypotheses. We are mainly interested in the following four measures: (1) RIX_i , RI at ICU discharge—this is a continuous variable with lower scores indicating poorer condition; (2) $I(RIX > RIE)_i$, whether RI at ICU discharge is greater than RI at ICU arrival—this is an indicator variable. Because the RI is a combination of many factors it is difficult to interpret the meaning of a change in RI of a few points. Therefore we also look at the probability of an improvement in the RI because it is easier to interpret as an improvement in the health status of a patient; (3) $\log(ICULOS)_i$, length-of-stay (LOS) in the ICU—as is standard practice, we take the logarithm of LOS in order to account for the heavy tail in its distribution; and (4) $I(SD)_i$, an indicator variable equal to 1 if the patient’s next unit is a step-down unit instead of a general ward or hospital discharge.

We use the following model:

$$Y_i = \alpha Workload_i + \lambda FloorBusy_i + \beta Z_i + \epsilon_i, \quad (1)$$

where Y_i is the outcome of interest for patient i and $Workload_i$ is a categorical variable as defined in Figure 2. We apply this model separately to each ICU. In addition to $Workload_i$ that represents ICU workload, we include the variable $FloorBusy_i$ to control for the busyness of downstream units, because the availability of downstream beds could affect the care provided in the ICU (Long and Mathews 2015). To compute $FloorBusy_i$, we pool all of the beds in step-down units and general medical-surgical wards the MICU and the SICU patients, respectively, are usually transferred to and calculate their average occupancy during the 12 hours prior to patient i ’s discharge from the ICU. We then define $FloorBusy_i = 1$ if the average is above the 75th percentile of its distribution and $FloorBusy_i = 0$ otherwise. Z_i is a vector of control variables that include (1) patient characteristics: RI at ICU arrival, age, gender, 27 DRG group dummies, 30 Elixhauser comorbidity dummies, and 16 dummies for principal procedure classification if the principal procedure is done before ICU discharge day; (2) patient type descriptions: whether an emergency admission, payor type dummies, and dummies for the unit before ICU; and (3) seasonality dummies including ICU admission month, ICU admission time of day, and ICU discharge day of week. Lastly, ϵ_i is the error term that captures all other factors that influence Y_i other than the regressors $Workload_i$, $FloorBusy_i$, and Z_i . Table 11 in the appendix provides a comprehensive list of these control variables as well as their summary statistics.

We estimate (1) with linear regression when Y_i is a continuous variable—when $Y_i = RIX_i$ and $Y_i = \log(ICULOS)_i$. When the outcome is a binary variable ($Y_i = I(RIX > RIE)_i$ and $Y_i = I(SD)_i$), we estimate (1) with logit regression. That is, the following specification is used:

$$\text{logit}(Pr(Y_i)) = \ln\left(\frac{Pr(Y_i)}{1 - Pr(Y_i)}\right) = \alpha Workload_i + \lambda FloorBusy_i + \beta Z_i. \quad (2)$$

4.3 Is RIX a meaningful measure?

Before we move on to the results, we first compare the RI at ICU discharge (RIX) to proxies of patient health status at ICU discharge and longer-term patient outcomes that are traditionally considered in the literature.

How do we verify if RIX or RI is at all meaningful? Accurately determining if an RI is “correct” is not possible, and we can only check its relationship to other outcome measures. Despite the fact the RI is a relatively new score (Rothman et al. 2013), there have been quite a number of studies that showed its strong association to patient outcomes measures, including mortality rates, ICU and hospital readmission rates, and other acuity scores—see details in Section 3.1. To illustrate these associations for our data set, in Table 3 we show the relationship between RIX and three patient outcome measures: ICU readmission ($reICU$), 30-day hospital readmission ($reHosp$), and in-hospital mortality ($death$). The MICU and the SICU patients, respectively, are split uniformly across groups 1 to 4 such that group 1 contains the most severe patients (i.e., patients with the lowest RIX) and group 4 contains the least severe patients (i.e., patients with the highest RIX). As expected, the patients in group 1 tend to have the highest ICU readmission, 30-day hospital readmission, and in-hospital mortality rates, and the patients in group 4 tend to have the lowest rates.

Table 3: RI at ICU discharge (RIX) Compared to RI at ICU arrival (RIE) and Longer-term Patient Outcomes

RIX	MICU							SICU						
	RIE				$reICU$	$reHosp$	$death$	RIE				$reICU$	$reHosp$	$death$
	Gr 1	Gr 2	Gr 3	Gr 4				Gr 1	Gr 2	Gr 3	Gr 4			
Gr 1	56.1%	29.9%	10.7%	3.4%	9.6%	30.9%	13.8%	45.0%	33.5%	13.1%	8.5%	12.4%	22.7%	5.6%
Gr 2	26.5%	35.9%	26.7%	11.0%	7.7%	28.2%	6.3%	31.4%	31.8%	23.3%	13.6%	5.9%	14.0%	0.8%
Gr 3	13.2%	23.7%	34.9%	28.2%	4.6%	25.2%	2.0%	14.1%	24.4%	35.5%	26.0%	2.4%	17.6%	0.8%
Gr 4	4.5%	10.3%	27.9%	57.3%	4.4%	24.7%	0.0%	9.4%	11.7%	27.0%	52.0%	3.2%	12.5%	0.0%
Total	25.1%	25.0%	25.0%	25.0%	6.6%	27.3%	5.5%	25.0%	25.4%	24.7%	24.9%	6.1%	16.7%	1.8%

Note. The MICU and the SICU patients, respectively, are split uniformly across groups 1 to 4 such that group 1 contains the most severe patients (i.e., patients with the lowest RIX) and group 4 contains the least severe patients (i.e., patients with the highest RIX). The 25th, 50th, and 75th percentile values of RIX among the MICU patients were 38.2, 52.8, 68.9, and 46.7, 57.6, 69.5 among the SICU patients.

The next question then is whether RIX has added value compared to longer-term patient outcome measures such as readmission and mortality rates. We believe that RIX as an intermediate outcome before the more common outcome measures has value because it provides a more accurate depiction of the state of the system *at the time* of the discharge event. Downstream outcome measures such as readmission and mortality rates are affected by operational factors in the hospital after ICU discharge. For example, suppose congestion in the ICU leads a patient to leave the ICU sicker. This means that downstream units have to care for this sicker patient and are stressed. In some cases they may be unsuccessful, but in the cases when they are successful at keeping the patient from dying or being readmitted to ICU, we might conclude that congestion had no impact.

Another benefit of the RI scores is its ability to track the changes in patient condition, which can help in applying targeted interventions (e.g., for the most severe group at ICU discharge) more effectively. KC and Terwiesch (2012) divided their cardiac surgery ICU patients into four groups by their initial acuity score to quantify the effect of congestion on LOS for each group (see Table 7 of Kc and Terwiesch (2012)). An important assumption they make is that the patients who leave the ICU are of the same relative acuity as when they came, as mentioned on page 61 of Kc and Terwiesch (2012): “We should note that the acuity level is based on the presurgery condition and diagnosis of the patient, not their medical condition at the time of discharge from the ICU. Given that the measures of acuity (age, gender, and procedure type) are fixed during a hospital stay, and given that we do not observe time-varying health status of patients, our acuity measure is fixed for a given patient’s stay.”


Because in our data we observe the time-varying health status of patients, we are able to compare RIX to RI at ICU arrival (RIE) in Table 3. Similar to how the RIX groups are formed, we uniformly split the MICU and the SICU patients into four groups depending on their RIE . We observe that while the transition from the RIE groups to the RIX groups is certainly not random, the two groups do not coincide either. This suggests that devising discharge strategies that depend on the initial acuity measure (e.g., the one suggested by Kc and Terwiesch (2012)) might not be fully effective, and that it is important to be able to observe time-varying health status of patients in order to develop successful decision support systems.

Previous studies have used LOS as a proxy for patient health status at discharge by assuming that all things being equal a patient who spends longer in the ICU will leave healthier than one who spends less time. The fact that two patients of similar relative health status can end up at different statuses, as shown in Table 3, indicates that LOS is of limited value as a measure of health status.

4.4 Impact of Workload for the Medical ICU Patients

In this subsection, the unit of analysis is individual MICU patient. We first compare the distributions of RIX by workload states: see Figure 3. If workload does not have any impact on ICU discharge practices, the RIX cumulative probability functions (cdfs) of the different workload states will look alike. Whereas the top left graph suggests that the RIX distribution does not depend on census (i.e., low-census versus high-census), the top right graph suggests that more acute patients are discharged from high-census/high-acuity compared to high-census/low-acuity. (The p-value for the two-sample Kolmogorov-Smirnov (KS) test for equality of distribution functions is 0.027—the distributions of RIX are statistically different at the 0.10 significance level.) In other words, the top right graph suggests that patients discharged when others are more acute tend to be more acute. For instance, the probability that RIX is smaller than or equal to 40 is 0.24 in

Figure 3: CDFs of RI at ICU discharge (RIX) for the Medical ICU Patients (I : low-census, II : high-census/low-acuity, and III : high-census/high-acuity)



high-census/low-acuity, and the same probability is 0.30 in high-census/high-acuity.

The three graphs on the bottom row of Figure 3 show that the RIX distribution depends on patients' next unit as well. After ICU stays, depending on the condition of the patients and the care path, patients can be either transferred to the step-down units, transferred to the general medical-surgical wards, or discharged from the hospital.⁴ Because step-down units provide a higher level of care compared to general wards, the least medically stable patients are transferred to a step-down unit. Hence, we expect the RI of patients going to the step-down units to be lower, as we see in the bottom left graph of Figure 3.

More interestingly, the bottom middle graph of Figure 3 suggests that census affects RIX even among the patients who go to a step-down unit. That is, more acute patients are sent to step-down units when the census is high (The p-value for the KS test is 0.029—the RIX distribution is statistically different at the 0.10 significance level). Lastly, the bottom right graph of Figure 3 further compares high-census/low-acuity and high-census/high-acuity for only the patients who are transferred to the step-down units. Note that the cdfs are not as smooth as the others because of the small sample sizes in each group—for instance, there are 18 patients in the “SD & high-census/low-acuity” group. We find that more acute patients are sent to step-down units from high-census/high-acuity compared to low-census (p-value 0.061), but there is no statistical difference in

⁴Note that we excluded ICU visits whose next units are the operating room, interventional radiology, catheterization and electrophysiology (Cath/EP), hemodialysis, or in-unit death because deciding when to send a patient to these next units are unlikely to be affected by the ICU workload; see Section 3.3.

the RIX distributions in high-census/low-acuity compared to low-census at the 0.10 significance level.

Table 4: Effect of Workload for the Medical ICU Patients

Workload	RIX			$I(RIX > RIE)$			$\log(ICULOS)$			$I(SD)$		
Low-census	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)
High-census	-0.02 (0.83)			-0.03 (0.14)			-0.15*** (0.04)			0.50* (0.21)		
High-census/low-acuity		1.64 (1.14)	-0.33 (1.45)		0.17 (0.20)	-0.19 (0.26)		-0.13** (0.05)	-0.06 (0.07)		0.10 (0.31)	0.03 (0.37)
High-census/high-acuity		-1.45 (1.07)	-2.43* (1.15)		-0.19 (0.17)	-0.33+ (0.18)		-0.17*** (0.05)	-0.14** (0.05)		0.77** (0.26)	0.80** (0.28)
FloorBusy	0.66 (0.87)	0.40 (0.88)	-0.87 (1.00)	0.09 (0.15)	0.06 (0.15)	-0.16 (0.17)	0.06* (0.04)	0.06* (0.04)	0.10* (0.05)	-0.91** (0.28)	-0.85** (0.29)	-0.84* (0.35)
High-census/low-acuity x FloorBusy			4.87* (2.17)			0.88* (0.40)			-0.16 (0.10)			0.19 (0.67)
High-census/high-acuity x FloorBusy			5.49+ (2.95)			0.83+ (0.45)			-0.12 (0.14)			-0.34 (0.86)
Observations (Pseudo) R^2	1864 0.528	1864 0.529	1864 0.531	1856 0.180	1856 0.181	1856 0.184	1864 0.312	1864 0.312	1864 0.313	1835 0.183	1835 0.186	1835 0.187

Note. Robust standard errors in parentheses. + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Next, we use the econometric model described in Section 4.2 to examine the impact of workload on RIX as well as other outcomes defined in Section 4.2. We control for other factors including patient characteristics and seasonality factors. In Table 4, we show the results from three different models for each outcome. Comparing the first model to the second model (the first column of each outcome to the second column of each outcome), we observe that examining high-census/low-acuity and high-census/high-acuity separately is important. For instance, our finding of shortened ICU LOS in a congested MICU agrees with the findings of [Kc and Terwiesch \(2012\)](#) in a cardiac ICU. In addition, using our new workload measure, we find that the ICU LOS is even more reduced in high-census/high-acuity, but less so in high-census/low-acuity. This pattern applies to most of the results in Table 4. That is, if we used the traditional workload measure only (high-census), we would have observed either no statistically significant impact of workload (on RIX) or smaller impact of workload (on $\log(ICULOS)$ and $I(SD)$). The results using our workload measure reveal that high-census/high-acuity is driving much of the impact of workload. We note that there are patients who are stabilized at a lower than average RI, and who would have been kept longer in the unit for observation purpose in a less congested ICU (low-census), without necessarily an increase in RI. But in a crowded ICU (high-census/high-acuity), this observation time could be shortened without an impact on RIX . Such patients will dampen the observed effect of workload on RIX , especially if those patients are more prevalent in high-census/high-acuity, which is likely the case.

In the third columns of each outcome, we examine the interaction between ICU workload and floor cen-

sus. We observe that including the interaction effects strengthens the effect of high-census/high-acuity. The coefficients of the interaction terms in the *RIX* models show that busy floor, together with high workload in the ICU, leads ICUs to discharge less acute patients. This suggests that when there is no room on the floor, the impact of ICU workload is mitigated because the MICU cannot discharge as many patients as it would have if there were more rooms available downstream. This is consistent with the assumptions underlying our simulation model. We can imagine that the patients are ordered from highest to lowest RI and the physicians discharge the patients from top to bottom in this list. If the ICU is crowded physicians need to make room by going deeper into this list and thus patients with lower RI values. If the downstream unit is crowded there is a limit to how deep into the list physicians can go and thus the average RI of a discharged patient is higher than if they could go deeper into the list. In addition, we observe that congestion in downstream units is associated with longer ICU LOS and lower odds of being transferred to a SD.


We also note that whereas we cannot, due to a lack of data, separate the ICU LOS into service time and time-to-transfer as Long and Mathews (2015) did, our results provide evidence that shortening ICU LOS is not just a matter of shortening the time-to-transfer, contrary to Long and Mathews (2015). Our results indicate that high workload is associated with shorter ICU LOS and also lower *RIX*. If only the time-to-transfer is shortened, then it should not have any health effect. However, our evidence of lower *RIX* in high-census/high-acuity indicates that there indeed is some health effect.

In summary, we find that in high-census/high-acuity (compared to low-census), on average, *RIX* drops by 2.43 points (a 5% decrease), the chance of having improved RI upon ICU discharge decrease by 5.9% (a 9% decrease), ICU LOS drops by 9.3 hours (a 13% decrease), and the chance of being transferred to a step-down bed increases by 7.5% (a 8% increase).

4.5 Impact of Workload for the Surgical ICU Patients

In this subsection, we repeat the analyses of Section 4.4 for the SICU patients. Figure 4 compares the SICU patients' *RIX* distributions by workload states. The two graphs on the top row suggest that the *RIX* distribution depend on the census (p-value 0.055—the distributions of *RIX* are statistically different at the 0.10 significance level), and that this difference is driven by high-census/high-acuity (p-value 0.001—between low-census and high-census/high-acuity). As expected, the bottom left graph of Figure 4 shows that the RI of patients going to the step-down units tend to be lower. The bottom middle graph suggests that more acute patients are sent to step-down units when the census is high, though we did not find statistical difference at the 0.10 significance level. Furthermore, more acute patients are sent to step-down units from high-census/high-acuity compared to low-census (p-value 0.033), but there is no statistical difference in the *RIX* distributions

Figure 4: CDFs of RI at ICU discharge (RIX) for the Surgical ICU Patients (I : low-census, II : high-census/low-acuity, and III : high-census/high-acuity)



in high-census/low-acuity compared to low-census at the 0.10 significance level.

Table 5: Effect of Workload for the Surgical ICU Patients

Workload	RIX			$I(RIX > RIE)$			$\log(ICULOS)$			$I(SD)$		
Low-census	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)
High-census	-3.18**			-0.62**			-0.06			0.86**		
	(1.18)			(0.23)			(0.06)			(0.29)		
High-census/low-acuity		-2.03	-0.47		-0.37	-0.41		-0.01	-0.03		0.27	0.40
		(1.34)	(1.64)		(0.26)	(0.34)		(0.07)	(0.09)		(0.37)	(0.40)
High-census/high-acuity		-5.34**	-4.15*		-1.07**	-1.07**		-0.11	-0.11		1.68***	1.68***
		(1.79)	(2.06)		(0.36)	(0.40)		(0.40)	(0.11)		(0.40)	(0.42)
FloorBusy	2.63*	2.52*	3.77**	0.28	0.26	-0.24	0.02	0.01	0.02	-0.71*	-0.64*	-0.57+
	(1.05)	(1.05)	(1.23)	(0.20)	(0.21)	(0.25)	(0.05)	(0.05)	(0.06)	(0.29)	(0.29)	(0.33)
High-census/low-acuity x FloorBusy			-4.18			0.10			0.03			-0.42
			(2.67)			(0.53)			(0.15)			(0.90)
High-census/high-acuity x FloorBusy			-4.53			0.04			-0.13			0.00
			(3.29)			(0.61)			(0.15)			(0.61)
Observations	1036	1036	1036	1029	1029	1029	1036	1036	1036	950	950	950
(Pseudo) R^2	0.428	0.430	0.432	0.305	0.308	0.308	0.382	0.383	0.384	0.328	0.337	0.338

Note. Robust standard errors in parentheses. + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

As we have seen for the MICU patients, Table 5 shows that high-census/high-acuity is driving much of the impact of workload. When we include the downstream census interaction terms with the SICU workload, we find that the impact on patient health status of being discharged (RIX) in high-census/high-acuity is weakened. This is different than what we saw for the MICU patients in Table 4. The MICU and the SICU

are different in many ways, which could play a role in explaining the different effect of the downstream unit. Some key differences are the patients in the two units (e.g., medical versus surgical patients), the sizes of the ICUs (the MICU has 36 staffed beds whereas the SICU has 21), and the approaches to using step down units. MICU patients tend to be sicker and more complex cases than SICU patients who tend to have a better defined care path. RIs for MICU patients have greater variance, both at arrival to the unit and at departure (see Table 2). From queueing theory we know that system size impacts performance under heavy workloads. So the fact that the MICU is significantly larger than the SICU means that we should not expect identical performance across the two ICUs. The MICU is more of a closed ICU while the SICU is more open⁵, which could lead to different interactions with step-down units and different access to step-down beds (e.g., in the hospital we study, there are 15 step-down beds where the MICU patients are usually sent to and 23 step-down beds for the SICU patients).

In summary, we find that in high-census/high-acuity (compared to low-census), on average, *RIX* drops by 4.15 points (a 7% decrease), the chance of having improved RI upon ICU discharge decrease by 16.8% (a 26% decrease), ICU LOS drops by 7.8 hours (a 12% decrease), and the chance of being transferred to a step-down bed increases by 21.8% (a 117% increase).

5 Impact of Patient Health Status at ICU Discharge on Longer-term Patient Outcomes

In the previous sections we have established that when the ICUs we studied were operating with high workload (high-census/high-acuity), discharged patients tended to be sicker. In addition to the simple analysis in Section 4.3, we quantify the impact of *RIX* on various longer-term patient outcome measures, to help put the impact of lower *RIX* into perspective. We consider six longer-term patient outcomes: RI at hospital discharge ($RIHX_i$), hospital LOS after the ICU stay ($\log(LOS_{afterICU})_i$), ICU readmission during the same

⁵The MICU is a “closed” ICU whereas the SICU is an “open” ICU in the hospital we study. An ICU is “closed” if all care in the ICU is directed by intensivists. An ICU is “open” if the intensivists are consultants without primary responsibility for the patient and separate attending physicians direct the patient care. (In the medical literature, studies have examined the advantages and drawbacks of having a closed ICU versus an open ICU: see [Brilli et al. \(2001\)](#), [Pronovost et al. \(2002\)](#), and [Pronovost et al. \(2006\)](#).) This means that in the MICU we study, only the MICU intensivists are treating the patients and when the patients leave the MICU, the MICU intensivists hand-off the care to other physicians. We believe that this leads the MICU intensivists to be less flexible about discharging a patient to another unit if the patients have not achieved a sufficient health status and thus ceteris paribus they discharge less acute patients. In the SICU we study the patients’ attending physicians are their surgeons, and the surgeons follow the patients throughout their stay in the hospital. As a result the surgeons know that they can keep an eye on the patients even after the patients leave the SICU. We suggest that this leads them to be more flexible about discharging a patient from the SICU.

hospitalization ($I(reICU)_i$), hospital readmission within 30 days from hospital discharge ($I(reHosp)_i$), in-hospital death during the same hospitalization ($I(death)_i$), and total hospital cost ($\log(Cost)_i$). Because patients who are directly discharged from the ICU will have $RIHX_i = RIX_i$, $LOSafterICU_i = 0$, $I(reICU)_i = 0$ and $I(death)_i = 0$, we study these four patient outcome measures over the subset of patients whose subsequent LOS after ICU is not zero (i.e., $LOSafterICU_i > 0$). Note that all of these variables are computed only over the patients who did not have in-hospital death, except for $I(death)_i$. Table 2 provides summary statistics of the longer-term patient outcome measures.

As was done in Section 4.2, we use linear regression when the outcome variable is continuous— $RIHX_i$, $\log(LOSafterICU)_i$, and $\log(Cost)_i$ —and logit regression when the outcome variable is binary— $I(reICU)_i$, $I(reHosp)_i$ and $I(death)_i$.⁶

Table 6 shows that higher RI at ICU discharge (RIX) is associated with better longer-term patient outcomes. The table also shows the average predicted values (APV) at RIX and at $RIX + 5$. That is, we assume that all patients are discharged with 5 more RI points from the ICU; 5 more RI points is chosen because Table 5 shows that the SICU patients' RIX decrease by about 5 points on average when discharged from high-census/high-acuity. We do not present the in-hospital mortality result for the SICU patients; because the in-hospital mortality rate among the SICU patients was low (1.9%), the majority of the observations were dropped during regression analysis due to some variables being perfect predictors of in-hospital mortality. The results in Table 6 support our claim that hospitals should avoid being in high workload states (measured by both census and patient acuity), because the impact of having lower RIX upon discharge due to high workload on longer-term patient outcomes is not negligible.

Table 6: Effect of RI at ICU discharge (RIX) on Longer-term Patient Outcomes

	MICU						SICU				
	$RIHX$	$\log(LOS)$	$I(reICU)$ after ICU)	$I(reHosp)$	I(in-hospital mortality)	$\log(Cost)$ mortality)	$RIHX$	$\log(LOS)$ after ICU)	$I(reICU)$	$I(reHosp)$	$\log(Cost)$
RIX	0.34*** (0.02)	-0.01*** (0.00)	-0.01** (0.00)	-0.00 (0.00)	-0.03*** (0.00)	-0.01*** (0.00)	0.33*** (0.03)	-0.02*** (0.00)	-0.02** (0.01)	-0.01** (0.00)	-0.01*** (0.00)
Observations	1796	1798	1628	1855	1825	1864	978	979	885	1004	1036
(Pseudo) R^2	0.59	0.33	0.24	0.07	0.36	0.43	0.48	0.45	0.43	0.18	0.51
APV at RIX	69.93	5.04 days	7.19%	27.39%	6.13%	\$ 24.4K	76.03	4.91 days	6.52%	16.92%	\$ 36.1K
APV at $RIX + 5$	71.64	4.68 days	6.65%	27.12%	5.01%	\$ 23.4K	77.69	4.46 days	5.71%	15.68%	\$ 34.2K

Note. APV stands for Average Predicted Value. Robust standard errors in parentheses. ⁺ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

⁶We have tested whether the RI at ICU discharge (RIX) is endogenous. That is, are there any unobserved factors related to patients clinical acuity that are affecting both RIX and downstream outcomes? Such unobservables will be positively correlated with lower RIX and adverse downstream outcomes, generating biases in the estimates of the causal effect of RIX . Using ICU workload as the instrument for RIX (similar to Kc and Terwiesch (2012)), we failed to reject the null hypothesis that RIX is exogenous. After all, RIX is an objective and composite measure of patient acuity computed based on observed data.

6 Interventions

In this section, we explore what an ICU manager could do to reduce the instances of high-census/high-acuity or to reduce their negative impact. In particular we will focus on reducing seasonality in the patient arrival volume and patient mix.

[Kc and Terwiesch \(2012\)](#) note, “When the ICU is full, the decision maker is confronted with the dilemma of whether to discharge an existing patient early or to cancel the surgery for a scheduled patient (because this patient would immediately require an ICU bed).” While this statement is true it does not give a manager much direction on how to improve the condition. At a higher level there are two approaches for managers to pursue. First, they can make interventions that mitigate the negative impact of a premature discharge, and second they can make interventions that reduce the likelihood of making premature discharges.

Here we briefly describe two possible ways to mitigate the negative impact of premature discharges. One approach is to develop predictive models that can help physicians better identify the lowest risk patients to discharge when an ICU is crowded. Another approach is to improve the quality of care in downstream units so that a patient prematurely discharged from an ICU is more likely to recover without a need to return to the ICU. This may require additional staff training or the development of new procedures for the higher risk patients. All these approaches are beyond the scope of the current paper.

We instead focus on reducing the likelihood of being in a state requiring premature discharges. Of course in theory the ICU can be expanded but this is a long-term solution that will also, in practice, lead to more demand to justify the cost and thus the core problem will remain. Canceling surgeries when an ICU is crowded is an immediate possible intervention. But as [Kc and Terwiesch \(2012\)](#) report, this happens infrequently because hospitals, surgeons and patients are loathe to do it.

We know from the principles of operations research that variability creates delays and wasted capacity in service processes; this idea has been applied to improve hospital operations, e.g., see [Bekker and Koeleman \(2011\)](#), [Chow et al. \(2011\)](#), [Helm and Van Oyen \(2014\)](#). One source of variability is in the patient arrival to the ICU. In [Figure 5](#) we show the arrival rates of patients by day-of-week, average RI of arriving patients, and the average daily census in both the MICU and SICU. In [Figure 6](#) we break out the SICU arrivals into those coming from the ED and those coming from the OR. We can see that for the SICU there is clearly day-of-week seasonality in both the number of arriving patients and the average RI. The seasonality is primarily among the patients coming from the OR, who are mainly scheduled for elective surgery. These surgeries could be scheduled more uniformly across the days of the week leading to a less variable SICU arrivals. A smoother arrival process should lead to the ICU being in high-census/high-acuity less often and then to fewer prematurely discharged patients.

To investigate the potential impact of smoothing the surgery schedule on premature discharges we use the simulation model from Section 4.1 for the SICU. In the simulation model, we assume the arrival rate and the mean *RIE* to be the same across different days of the week, and set the arrival rate to be 5 patients per day and the mean *RIE* equal to 52. This model's performance is presented in the first column of Table 7. In the three additional sets of simulations, we introduce seasonality to the arrival rates and to the daily mean *RIE*. To ensure that the arriving workload is the same, we fix the weekly arrival rate to be 35 patients across all four sets of simulations, but use the percentage of arrivals in the leftmost graph of Figure 5 to make the daily arrival rate change. For the second set of simulation (second column of Table 7), we introduce the seasonality observed in the data (i.e., as shown in the middle graph of Figure 5) to the daily mean *RIE*. For the third set of simulation (third column), we rearrange the daily mean *RIE* so that when there are many arrivals, healthier patients (higher *RIE*) arrive, so that the net effect of seasonality is moderated. For the fourth set of simulation (fourth column), we rearrange the daily mean *RIE* so that when there are many arrivals, sicker patients (lower *RIE*) arrive, so that the net effect of seasonality is exacerbated.

Table 7: Results from simulations with varying degree of seasonality in the arrival volume and patient mix. Average and its 95% confidence interval over 1,000 iterations.

	No seasonality	As observed in data	In Favor	Against
Simulation Parameters				
Arrival rate	[5,5,5,5,5,5]	[3,6,5,2,6,6,5,2,5,2,6,1,3,1]	[3,6,5,2,6,6,5,2,5,2,6,1,3,1]	[3,6,5,2,6,6,5,2,5,2,6,1,3,1]
Mean <i>RIE</i>	[52,52,52,52,52,52]	[51,0,52,5,55,1,54,0,49,7,49,3,49,9]	[49,7,52,5,55,1,51,0,49,9,54,0,49,3]	[54,0,49,9,49,3,51,0,52,5,49,7,55,1]
Simulation Results				
Mean <i>RIX</i>	64.5±0.01	64.5±0.01	64.5±0.01	64.2±0.01
Avg occupancy (% high-census/high-acuity)	15.6±0.02 (8.6±0.1%)	15.8±0.02 (10.6±0.1%)	15.5±0.02 (9.0±0.1%)	16.3±0.02 (17.0±0.2%)
Sunday	15.6±0.03 (8.6±0.2%)	15.4±0.03 (12.3±0.3%)	14.7±0.03 (6.4±0.2%)	16.2±0.03 (9.0±0.2%)
Monday	15.5±0.03 (8.5±0.2%)	14.7±0.03 (8.1±0.2%)	13.7±0.03 (5.1±0.2%)	14.4±0.03 (3.0±0.1%)
Tuesday	15.6±0.03 (8.6±0.3%)	15.6±0.03 (9.7±0.3%)	14.9±0.03 (7.6±0.2%)	14.5±0.03 (7.2±0.2%)
Wednesday	15.6±0.03 (8.7±0.2%)	17.0±0.03 (7.9±0.2%)	16.6±0.03 (7.2±0.2%)	16.5±0.03 (23.0±0.4%)
Thursday	15.6±0.03 (8.8±0.2%)	16.0±0.03 (5.1±0.2%)	15.8±0.03 (9.5±0.2%)	17.2±0.03 (25.8±0.4%)
Friday	15.6±0.03 (8.5±0.2%)	15.3±0.03 (9.4±0.3%)	15.8±0.03 (13.8±0.3%)	17.5±0.03 (20.5±0.3%)
Saturday	15.6±0.03 (8.5±0.2%)	16.5±0.03 (21.9±0.3%)	17.1±0.03 (13.6±0.3%)	17.9±0.03 (30.4±0.4%)
% premature discharges	3.43±0.07%	4.26±0.07%	4.10±0.07%	6.74±0.09%

Table 7 suggest that by smoothing the arrivals and the mean *RIE* (comparing second column to first column), the percentage of time in high-census/high-acuity could be reduced, from 10.6% to 8.6%. As a result, the percentage of premature discharges could drop by 0.83% from 4.26% to 3.43%, which is a 19% decrease. The arrival rate and mean *RIE* observed in the data reveal that the SICU in our study is already aligning high patient arrival days with healthier patients to some extent. The results in column 3 and column 4 suggest that in case seasonality is unavoidable, it is important to align high patient arrival days with healthier patients. For instance, the simulation results show that if the SICU we study instead admitted sicker patients when there are many arrivals (column 4), the ICU would be in high-census/high-acuity 17.0% of the time instead of 10.6%, which would lead to a 58% increase in the percentage of premature discharges, from 4.26%

to 6.74%.

7 Conclusion

In this study, we have examined the impact of ICU workload on health status and care path of patients discharged from ICUs. In doing so, we used a detailed patient dataset of two ICUs at a major U.S. teaching hospital. Previous studies have examined the impact of ICU workload only using proxies of health status of patients or longer-term patient outcomes due to the absence of a dynamic measure of patient acuity. The Rothman Index (RI) allows us to examine the direct relationship between workload and health status of patients in this study. We found that more acute patients are discharged when the ICU has high workload. Furthermore, leveraging the RI, we measured the workload using not only the census (as is traditionally done) but also the acuity of the patients currently in the ICU. By studying the impact of this acuity-adjusted workload on the health status and the care path of patients discharged from the ICUs, we show that the acuity mix of patients, in addition to the census, drives the changes in patient flow. We also showed that poorer health status at ICU discharge was associated with worse longer-term patient outcomes, including longer post-ICU LOS, higher in-hospital mortality rate and higher total hospital costs.

At its heart patient flow in a hospital is about which patients get which resources and when. Given the high utilization and expense of intensive care units there will be continued interest in understanding patient flow through ICUs. Our study provides insights into how researchers should proceed in studying ICU patient flow. We demonstrated that (1) census is an incomplete measure of workload, and hence ICUs need to track the changes in patient acuity in addition to census in order to use such information to prevent reaching high occupancy with many high acuity patients, and (2) future studies of ICU workload should take patient acuity into account in measuring the workload because this study shows that patient acuity in combination with census are affecting who is discharged from the ICU and when.

We used a simulation study to explore how smoothing arrivals can alleviate heavy ICU workloads that impact patient health status. We found that when possible smoothing arrivals (and thus workload) can significantly reduce instances of high-census/high-acuity, which significantly reduces the number of prematurely discharged patients.

This study also shows that the impact of workload in ICUs can differ in magnitude and in how ICU workload interacts with the workload of downstream units across ICUs within a single hospital. We speculate that these differences could be due to a number of factors including patient type (e.g., medical versus surgical patient), size of the ICU, different workplace culture (e.g., closed ICU versus open ICU), and access to downstream beds. Our findings suggest that a blanket statement on the impact of workload in ICUs, which

is often done in the operations management literature, could be misleading. Future work should seek to understand how different characteristics of ICUs interact with workload to influence patient care and health status.

Lastly, while our new measure of ICU workload is shown to have a great advantage over the traditional workload measure that takes into account only patient census, it is questionable whether it is the *best* workload measure one can use. As Halpern (2009) claimed, we need to continue looking for the most parsimonious measure of ICU workload that can capture the differences in care: “Development of an appropriately accurate and parsimonious measure of ICU capacity strain may augment the precision of future critical care outcomes research by reducing unexplained variance attributable to temporal fluctuations in ICU-level factors; elucidate organizational characteristics that make some ICUs better able to withstand high capacity strain without substantive degradations in quality; and enhance the transparency of critical care rationing while helping to improve its equity and efficiency, thereby promoting the ethics of this inevitable practice (page 648 of Halpern (2009)).”

References

- Allon, G, S Deo, W Lin. 2013. The impact of size and occupancy of hospital on the extent of ambulance diversion: Theory and evidence. *Operations Research* **61**(3) 544–562.
- Anderson, D, C Price, B Golden, W Jank, E Wasil. 2011. Examining the discharge practices of surgeons at a large medical center. *Health care management science* **14**(4) 338–347.
- Batt, R J, C Terwiesch. 2016. Early task initiation and other load-adaptive mechanisms in the emergency department. *Forthcoming in Management Science* .
- Bekker, René, Paulien M Koeleman. 2011. Scheduling admissions and reducing variability in bed demand. *Health care management science* **14**(3) 237–249.
- Beswick, Susan, Pamela D Hill, Mary Ann Anderson. 2010. Comparison of nurse workload approaches. *Journal of Nursing Management* **18**(5) 592–598.
- Bradley, E, O Yakusheva, L I Horwitz, H Sipsma, J Fletcher. 2013. Identifying patients at increased risk for unplanned readmission. *Medical care* **51**(9) 761–766.
- Brilli, R J, A Spevetz, R D Branson, G M Campbell, H Cohen, J F Dasta, M A Harvey, M A Kelley, K M Kelly, M I Rudis, et al. 2001. Critical care delivery in the intensive care unit: defining clinical roles and the best practice model. *Critical care medicine* **29**(10) 2007–2019.
- Chan, C W, V F Farias, N Bambos, G J Escobar. 2012. Optimizing intensive care unit discharge decisions with patient readmissions. *Operations research* **60**(6) 1323–1341.
- Chow, Vincent S, Martin L Puterman, Neda Salehirad, Wenhai Huang, Derek Atkins. 2011. Reducing surgical ward

- congestion through improved surgical scheduling and uncapacitated simulation. *Production and Operations Management* **20**(3) 418–430.
- Chrusch, C A, K P Olafson, P M McMillan, D E Roberts, P 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.
- Danesh, V, L Guerrier, E Jimenez. 2012. Proactive vs. reactive rapid response systems: Decreasing unplanned icu transfers. *Critical Care Medicine* **40**(12) 1–328.
- Dobson, G, H-H Lee, E Pinker. 2010. A model of ICU bumping. *Operations research* **58**(6) 1564–1576.
- Elixhauser, A, C Steiner, D R Harris, R M Coffey. 1998. Comorbidity measures for use with administrative data. *Medical care* **36**(1) 8–27.
- Euroscore. 2007. URL <http://www.euroscore.org/>.
- Finlay, G D, M J Rothman, R A Smith. 2014. Measuring the modified early warning score and the rothman index: advantages of utilizing the electronic medical record in an early warning system. *Journal of Hospital Medicine* **9**(2) 116–119.
- Freeman, M, N Savva, S Scholtes. 2016. Gatekeepers at work: An empirical analysis of a maternity unit. *Forthcoming in Management Science*.
- Green, L V, S Savin, N Savva. 2013. “nursevendor problem”: Personnel staffing in the presence of endogenous absenteeism. *Management Science* **59**(10) 2237–2256.
- Green, Linda V. 2002. How many hospital beds? *INQUIRY: The Journal of Health Care Organization, Provision, and Financing* **39**(4) 400–412.
- Halpern, N A. 2009. Can the costs of critical care be controlled? *Current opinion in critical care* **15**(6) 591–596.
- Hasija, S, E Pinker, R A Shumsky. 2010. Om practice-work expands to fill the time available: Capacity estimation and staffing under parkinson’s law. *Manufacturing & Service Operations Management* **12**(1) 1–18.
- Helm, Jonathan E, Mark P Van Oyen. 2014. Design and optimization methods for elective hospital admissions. *Operations Research* **62**(6) 1265–1282.
- Iwashyna, T J, A A Kramer, J M Kahn. 2009. Intensive care unit occupancy and patient outcomes. *Critical care medicine* **37**(5) 1545.
- Jaeker, J B, A L Tucker. 2016. Past the point of speeding up: The negative effects of workload saturation on efficiency and patient severity. *Forthcoming in Management Science*.
- Kc, D S, C Terwiesch. 2009. Impact of workload on service time and patient safety: An econometric analysis of hospital operations. *Management Science* **55**(9) 1486–1498.
- Kc, D S, C Terwiesch. 2012. An econometric analysis of patient flows in the cardiac intensive care unit. *Manufacturing & Service Operations Management* **14**(1) 50–65.
- Kim, S-H, C W Chan, M Olivares, G Escobar. 2015. ICU admission control: An empirical study of capacity allocation and its implication for patient outcomes. *Management Science* **61**(1) 19–38.

- Kuntz, L, R Mennicken, S Scholtes. 2015. Stress on the ward: Evidence of safety tipping points in hospitals. *Management Science* **61**(4) 754–771.
- Kuntz, L, S Sülz. 2013. Treatment speed and high load in the emergency department—does staff quality matter? *Health care management science* **16**(4) 366–376.
- Long, E F, K S Mathews. 2015. Patients without patience: An econometric analysis of waiting in the intensive care unit. *Working Paper, UCLA Anderson School of Management* .
- Miranda, D R, M Jegers. 2012. Monitoring costs in the icu: a search for a pertinent methodology. *Acta Anaesthesiologica Scandinavica* **56**(9) 1104–1113.
- Miranda, D R, R Nap, A de Rijk, W Schaufeli, G Iapichino, et al. 2003. Nursing activities score. *Critical care medicine* **31**(2) 374–382.
- Piper, G L, L J Kaplan, A A Maung, F Y Lui, K Barre, K A Davis. 2014. Using the rothman index to predict early unplanned surgical intensive care unit readmissions. *Journal of Trauma and Acute Care Surgery* **77**(1) 78–82.
- Powell, A, S Savin, N Savva. 2012. Physician workload and hospital reimbursement: Overworked physicians generate less revenue per patient. *Manufacturing & Service Operations Management* **14**(4) 512–528.
- Pronovost, P J, D C Angus, T Dorman, K A Robinson, T T Dremsizov, T L Young. 2002. Physician staffing patterns and clinical outcomes in critically ill patients: a systematic review. *Jama* **288**(17) 2151–2162.
- Pronovost, P J, C G Holzmüller, L Clattenburg, S Berenholtz, E A Martinez, J R Paz, D M Needham. 2006. Team care: beyond open and closed intensive care units. *Current opinion in critical care* **12**(6) 604–608.
- Robert, R, J Reignier, C Tournoux-Facon, T Boulain, O Lesieur, V Gissot, V Souday, M Hamrouni, C Chapon, J-P Gouello. 2012. Refusal of intensive care unit admission due to a full unit: impact on mortality. *American journal of respiratory and critical care medicine* **185**(10) 1081–1087.
- Rothman, M J, S I Rothman, J Beals. 2013. Development and validation of a continuous measure of patient condition using the electronic medical record. *Journal of biomedical informatics* **46**(5) 837–848.
- Shi, Pengyi, Mabel C Chou, JG Dai, Ding Ding, Joe Sim. 2015. Models and insights for hospital inpatient operations: Time-dependent ed boarding time. *Management Science* **62**(1) 1–28.
- Shmueli, A, C L Sprung, E H Kaplan. 2003. Optimizing admissions to an intensive care unit. *Health Care Management Science* **6**(3) 131–136.
- Singer, D E, P L Carr, A G Mulley, G E Thibault. 1983. Rationing intensive care—physician responses to a resource shortage. *The New England journal of medicine* **309**(19) 1155–1160.
- Strand, K, H Flaatten. 2008. Severity scoring in the icu: a review. *Acta Anaesthesiologica Scandinavica* **52**(4) 467–478.
- Tan, T F, S Netessine. 2014. When does the devil make work? an empirical study of the impact of workload on worker productivity. *Management Science* **60**(6) 1574–1593.
- Tarnow-Mordi, W O, C Hau, A Warden, A J Shearer. 2000. Hospital mortality in relation to staff workload: a 4-year study in an adult intensive-care unit. *The Lancet* **356**(9225) 185–189.

- Tepas, Joseph J, J M Rimar, A L Hsiao, M S Nussbaum. 2013. Automated analysis of electronic medical record data reflects the pathophysiology of operative complications. *Surgery* **154**(4) 918–926.
- Wagner, J, N B Gabler, S J Ratcliffe, S ES Brown, B L Strom, S D Halpern. 2013. Outcomes among patients discharged from busy intensive care units. *Annals of internal medicine* **159**(7) 447–455.
- Welton, J. 2007. Mandatory hospital nurse to patient staffing ratios: Time to take a different approach. *Online Journal of Issues in Nursing* **12**(3).

A Supplementary Materials for Section 4

In Section 4, we propose a simple model of the operation of the ICU to generate hypotheses of what we should expect to see in our data if the model is correct. The main operating premises of our model are that the patient health status is completely determined by the RI, that patients are discharged according to their RI which care providers can perfectly discern and that new arrivals are given priority over current occupants of the ICU. Here we provide a detailed description of the simulation model and results of numerical experiments using realistic parameter values.

For each day d a patient i stays in an ICU, the patient's RI RI_i changes by a random amount $RIC_{i,d}$ drawn from a normal distribution with mean μ_{RIC} and standard deviation σ_{RIC} . When RI_i exceeds a threshold *Healthy*, patient i is ready for discharge to another less intensive unit of the hospital. We will call such patients: healthy. The implication of these assumptions is that a patient's LOS is tightly linked to the patient's RI at arrival to the unit. At the start of the day (d) the care providers discharge all patients who are healthy. Then a random number of new patient arrivals, NEW_d , from a Poisson distribution with mean λ is determined. If NEW_d exceeds the number of available beds, then the care providers will discharge additional patients from the ICU who are not yet healthy. These additional (premature) discharges will be done in order of healthiest (highest RI) patients first. Each new arrival j has an initial RI RIE_j drawn from a normal distribution with mean μ_{RIE} and standard deviation σ_{RIE} .

We simulate the MICU and the SICU, separately, using the above framework with parameter values derived from our data: see Table 8. We choose μ_{RIE} close to the mean RIE of the MICU and SICU patients, respectively (see Table 2). To set the value for μ_{RIC} , we divide $(RIX - RIE)$ for each patient by her length-of-stay, and compute the average. We choose the values for λ , σ_{RIE} , σ_{RIC} , NEW_d and *Healthy* close to what we observe in the data (e.g., *Healthy* close to the mean of RIX), but adjust them to make the simulated ICU occupancy and LOS close to what we observe in the data. We note that we set the standard deviation values σ_{RIE} and σ_{RIC} to be both 4, which is smaller than observed in the data (in the data we observe $\sigma_{RIE} \approx 20$ and $\sigma_{RIC} \approx 11$), because large standard deviations will result in too many simulated LOS that are either too long or too short. We think smaller standard deviation values in our simulation models are justified because in reality, is it not only the RI scores but also other variables (e.g., diagnoses and procedures) that determine the LOS. The RI scores in our simulations are used as a composite of all those variables, and hence these other variables should help reduce some of the variations in the actual RI distribution we observe in the data.

For each parameter set, we run 1,000 replications where each replication is run for 1,000 days. To get rid of any initial effects, we use the last 2000 patients of each replication for our analysis. We assign workload states

Table 8: Parameters Selected for the Medical ICU (MICU) and Surgical ICU (SICU) Simulation Experiments

Parameter	Description	MICU	SICU
NumBeds	Number of ICU beds	36	21
New_d	Number of new patient arrivals in day d . Drawn from a Poisson distribution with mean λ	10	5
RIE_i	Patient i 's initial RI. Drawn from a Normal distribution with $(\mu_{RIE}, \sigma_{RIE})$	(46, 5)	(52, 5)
$RIC_{i,d}$	Patient i 's change in RI in day d . Drawn from a Normal distribution with $(\mu_{RIC}, \sigma_{RIC})$	(4, 4)	(4, 4)
$Healthy$	Patient i is health if RI_i exceeds this threshold	55	61

to the 2000 patients in each replication in the following way. First, we define a patient j to be *MoreAcute* if RI_j is less than or equal to 52 in the MICU and 57 in the SICU. At the beginning of the day patient i is discharged from the ICU (before any discharges happen), we count the number of patients in the ICU other than patient i and the percentage of the *MoreAcute* patients among them. We use a threshold of 35 patients and 18 patients for the MICU and the SICU, respectively, to divide into low-census and high-census. We then further divide high-census into high-census/low-acuity and high-census/high-acuity by whether the percentage of more acute patients is less than or equal to 55% for the MICU and 50% for the SICU. These thresholds are deliberately selected to assign approximately 75% patients to low-census, 12.5% to high-census/high-acuity, and 12.5% to high-census/low-acuity. However, because of the integrality of the occupancy, we have on average (over 1,000 replications) 66.6% of the MICU patients discharged from low-census, 19.4% from high-census/low-acuity, and 14.0% from high-census/high-acuity and 68.9% of the SICU patients discharged from low-census, 17.7% from high-census/low-acuity, and 13.5% from high-census/high-acuity.

Table 9: Regression Results from the Medical ICU (MICU) and Surgical ICU (SICU) Simulation Experiments. The medians of workload states' coefficients and their corresponding p-values (in parentheses) from 1000 replications are reported.

	MICU			SICU		
Workload	RIX	$\log(ICULOS)$	LOS	RIX	$\log(ICULOS)$	LOS
Low-census	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)
High-census	-0.64 (0.00)	-0.07 (0.00)	-0.06 (0.03)	-0.46 (0.00)	-0.05 (0.05)	-0.04 (0.14)
Low-census	(ref)	(ref)	(ref)	(ref)	(ref)	(ref)
High-census/low-acuity	-0.28 (0.09)	-0.05 (0.11)	-0.04 (0.22)	-0.17 (0.29)	-0.03 (0.28)	-0.03 (0.37)
High-census/high-acuity	-1.14 (0.00)	-0.10 (0.00)	-0.09 (0.02)	-0.81 (0.00)	-0.07 (0.03)	-0.06 (0.12)

To test the impact of workload on RIX , we run linear regressions with RIX as the dependent variable and workload and RIE as the independent variables. To test the impact of workload on patients' length-of-stay, LOS , we run Poisson regressions with LOS as the dependent variable and workload and RIE as the independent variables, because, given the design of the simulation, LOS is a count variable. (In Section 4.2, we use linear regressions for $\log(LOS)$ because LOS in our data is measured in seconds. For comparison, we also provide the results of linear regressions for $\log(LOS)$ below, but we note that these models fail the assumption of normally distributed errors.) Because we run 1,000 replications of the MICU or the SICU simulation, we have 1,000 different coefficients for workload states and 1,000 corresponding p-values. We present their median values in Table 9.

The results presented in Table 9 are consistent with our hypotheses 1, 2, and 3 in Section 4. The results are also consistent with the outcome of the econometric analysis in Section 5. We have conducted extensive simulation experiments with varying parameter values (e.g., varying μ_{RIE} , σ_{RIE} , μ_{RIC} and σ_{RIC} , NumBeds, and λ) and the results provide robust support for our hypotheses.

B Additional Tables and Figures

Table 10: Sample Selection for the Medical ICU (MICU) and Surgical ICU (SICU)

Sample	MICU			SICU		
	Obs	% prior	% initial	Obs	% prior	% initial
All ICU visits from 2/1/2013 to 1/31/2014	3227	NA	100.0	1556	NA	100.0
Excluding non-first-time ICU	2754	85.3	85.3	1372	88.2	88.2
Excluding those that died during the ICU stay	2475	89.9	76.7	1307	95.3	84.0
Excluding transfers from neighboring hospitals	2466	99.6	76.4	1306	99.9	83.9
Excluding those whose next unit is another ICU or the operating room	2349	95.3	72.8	1228	94.0	78.9
Excluding those who left against medical advice or unknown discharge destinations	2324	98.9	72.0	1222	99.5	78.5
Excluding ICU length of stays shorter than 12 hours or longer than 21 days	2176	93.6	67.4	1172	95.9	75.3
Excluding those preceded by hospital stays that are longer than 7 days	2063	94.8	63.9	1130	96.4	72.6
Excluding those with no Rothman Index or missing patient information	2052	99.5	63.6	1068	94.5	68.6
Excluding those with in-hospital mortality	1939	94.5	60.1	1049	98.2	67.4
Excluding those classified as comfort-measure-only before hospital discharge	1864	96.1	57.8	1036	98.8	66.6

Figure 5: Percentage of arrivals, mean RI at ICU arrival (RIE) and mean daily average occupancy for the Medical ICU (MICU) and Surgical ICU (SICU) by day-of-week




Figure 6: Percentage of arrivals and mean RI at ICU arrival (RIE) of patients coming from the operating room (OR) or the emergency department (ED) to the Surgical ICU by day-of-week




Table 11: Descriptions and Summary Statistics of the Control Variables (Means only)

Measure	Description and Coding	MICU (n=1864)	SICU (n=1036)
RI at ICU Arrival	Coded as piecewise linear spline variables with knots at 40 and 65	45.4	51.8
Age	Coded as piecewise linear spline variables with knots at 45 and 65	61.1	58.8
Gender	Females were coded 1 and males 0	0.48	0.42
Principal procedure classification	Included if the principal procedure is done before ICU discharge day. The classification (e.g., operations on the eye) is provided by the Clinical Classification Software (CCS) codes. See http://www.hcup-us.ahrq.gov/toolssoftware/ccs/ccs.jsp for details. Categorical variable to denote operations on (nervous system, endocrine system, eye, ear, nose/mouth/pharynx, respiratory system, cardiovascular system, hemic and lymphatic system, digestive system, urinary system, male genital organs, female genital organs, Obstetrical procedures, musculoskeletal system, integumentary system, miscellaneous, none)	(0.013, 0.000, 0.001, 0.000, 0.001, 0.036, 0.135, 0.001, 0.107, 0.069, 0.000, 0.001, 0.000, 0.009, 0.008, 0.356, 0.263)	(0.061, 0.020, 0.068, 0.029, 0.042, 0.040, 0.093, 0.015, 0.230, 0.050, 0.005, 0.019, 0.006, 0.135, 0.033, 0.109, 0.133)
Major Diagnostic Category (MDC)	Formed by categorizing the DRG codes into 27 mutually exclusive diagnosis areas. The first group is pre-MDC (surgical transplant) and the last group is invalid and ungroupable DRGs. See http://health.utah.gov/opa/IBIShelp/codes/MDC.htm for details.	(0.009, 0.023, 0.001, 0.007, 0.206, 0.101, 0.110, 0.046, 0.017, 0.004, 0.064, 0.046, 0.001, 0.003, 0.002, 0.000, 0.013, 0.010, 0.199, 0.000, 0.041, 0.082, 0.000, 0.003, 0.000, 0.008, 0.006)	(0.076, 0.189, 0.005, 0.023, 0.037, 0.048, 0.101, 0.109, 0.107, 0.013, 0.021, 0.039, 0.005, 0.017, 0.010, 0.000, 0.005, 0.018, 0.032, 0.001, 0.000, 0.037, 0.000, 0.003, 0.088, 0.002, 0.015)
ER admission	Coded 1 if the hospitalization started from the emergency room and 0 otherwise	0.85	0.52
Payer	Categorical variable to denote (Medicare, Managed Medicare, Medicaid, Managed Care—includes Blue Cross and Commercial Managed Care, Other)	(0.46, 0.10, 0.24, 0.18, 0.01)	(0.33, 0.10, 0.20, 0.34, 0.03)
Previous unit	The unit patient was in right before ICU. Categorical variable to denote (step-down unit, general ward unit, none (direct hospital admission to the ICU), operating room, emergency room, various tests)	(0.06, 0.19, 0.09, 0.01, 0.64, 0.01)	(0.04, 0.13, 0.04, 0.45, 0.32, 0.02)
ICU admission month	Categorical variable to denote (Feb 2013, Mar 2013, Apr 2013, May 2013, June 2013, July 2013, Aug 2013, Sept 2013, Oct 2013, Nov 2013, Dec 2013, Jan 2014)	(0.08, 0.09, 0.09, 0.08, 0.08, 0.09, 0.07, 0.08, 0.09, 0.09, 0.08, 0.08)	(0.07, 0.09, 0.09, 0.08, 0.08, 0.09, 0.10, 0.06, 0.07, 0.09, 0.09, 0.09)
ICU admission hour of day	Categorical variable to denote ([12am, 3am), [3am, 6am), [6am, 9am), [9am, 12pm), [12pm, 3pm), [3pm, 6pm), [6pm, 9pm), [9pm, 12am))	(0.12, 0.10, 0.06, 0.10, 0.13, 0.17, 0.18, 0.14)	(0.10, 0.08, 0.06, 0.07, 0.15, 0.21, 0.19, 0.15)
ICU discharge day-of-week	Categorical variable to denote (Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday)	(0.14, 0.13, 0.16, 0.14, 0.16, 0.15, 0.12)	(0.12, 0.11, 0.14, 0.17, 0.16, 0.16, 0.14)
Downstream bed occupancy	The average occupancy of downstream beds during the 12 hours prior to the focal patient's discharge from the ICU. The average was 266.8 and 150.8 for the MICU and SICU patients, respectively.	0.25	0.25

Note. Summary statistics of 30 Elixhauser comorbidity indexes, computed from the ICD-9 codes and the DRG, are not provided due to space limitations.