Optimizing Patient Placement in Normal Care Units: An Instrumental Causal Forest Approach Minimizing Mortality

Johannes Cordier*

Chair of Health Economics, Policy and Management, School of Medicine, University of St. Gallen
Erasmus School of Health Policy and Management and Erasmus School of Economics, Erasmus University Rotterdam

01.04.2025

Abstract

Effective patient placement in normal care units (NCUs) is essential for optimizing clinical outcomes and resource utilization. This study examines the impact of NCU placement on in-hospital mortality using administrative data from Swiss university hospitals. Employing an instrumental variable causal forest approach, we estimate heterogeneous treatment effects while addressing selection bias through daily NCU busyness as an instrument.

Our findings highlight a trade-off between NCU specialization and utilization. To address this, we propose a minimax regret policy framework that optimizes patient placement by minimizing worst-case regret. The policy reduces mortality, balances NCU busyness, and maintains welfare metrics without requiring additional hospital capacity.

This study demonstrates the potential of data-driven policies to improve patient outcomes and resource allocation in resource-constrained healthcare settings.

*Corresponding Author, johannes.cordier@unisg.ch

1 Introduction

The placement of patients in normal care units (NCUs) during hospitalization is crucial for clinical outcomes, length of stay, and costs (Sharma et al., 2022; Handel et al., 2018). NCUs provide specialized treatment and monitoring for patients who do not require intensive care, yet high NCU utilization has been linked to increased mortality and adverse events (Abir et al., 2020; Sharma et al., 2022; Castagna et al., 2022; Boyle et al., 2013; Schilling et al., 2010; Kuntz et al., 2015; Wise, 2015; Madsen et al., 2014), primarily due to overburdened staff reducing individualized care. Conversely, lower utilization enhances treatment effectiveness by allowing more personalized care.
Hospital occupancy plays a central role in patient outcomes. High occupancy rates correlate with increased adverse events (Boyle et al., 2013), higher mortality (Schilling et al., 2010; Madsen et al., 2014; Abir et al., 2020), and safety tipping points in hospitals (Kuntz et al., 2015). Effective bed management has been associated with reduced mortality (Wise, 2015), and the impact of occupancy on mortality was notably pronounced during the COVID-19 pandemic (Castagna et al., 2022).
Patient placement within hospitals further influences outcomes. Studies show that suboptimal placement negatively affects care quality (Handel et al., 2018; Lloyd et al., 2005; Stowell et al., 2013; Alameda and Suárez, 2009; Lepage et al., 2009), particularly in trauma and specialized care settings. Out-lying patients due to bed shortages leads to worse nursing care and increased risks (Stowell et al., 2013; Alameda and Suárez, 2009). Moreover, emergency department overcrowding elevates mortality risks (Sprivulis et al., 2006; Elsayed et al., 2005).
When NCUs reach capacity, patients may be placed in alternative units, a practice known as ”out-lying” (Lloyd et al., 2005; Stowell et al., 2013; Alameda and Suárez, 2009; Lepage et al., 2009). While studies have examined occupancy-driven mortality effects (Sprivulis et al., 2006; Elsayed et al., 2005), this study focuses on placement decisions conditional on NCU utilization.
We hypothesize that patients in underutilized NCUs may experience lower mortality and shorter hospital stays than those in heavily utilized units. However, specialized NCUs may improve outcomes for specific conditions, creating a trade-off between utilization and specialization. Furthermore, patient outcomes are interdependent, as placement decisions influence subsequent utilization. Traditional strategies based on diagnoses and risk factors overlook these dynamic effects (Lloyd et al., 2005; Stowell et al., 2013).
This study addresses the following questions:

1. How can a patient placement policy be designed to minimize mortality while balancing utilization and specialization?

2. What are the welfare effects of implementing such a policy?

To answer these, I analyze Swiss inpatient data (2012–2020) covering surgical and internal medicine NCUs. I focus on patients with primary ICD-10 diagnoses in ischemic heart diseases (I2, I3), cerebrovascular diseases (I6, I7), and various cancers (C1–C4, C7), ensuring analytical consistency.
We propose a two-stage framework integrating instrumental variable causal forests with a minimax regret policy. First, heterogeneous causal effects of NCU admission are estimated using an instrumental variable causal forest (Wager and Athey, 2018a; Athey and Imbens, 2016), leveraging exogenous variation in utilization to mitigate confounding. Patient-specific policy scores quantify expected benefits of NCU admission. Second, a minimax regret criterion is employed to optimize placement decisions dynamically, adjusting for utilization effects.
This study makes two contributions. First, it introduces a methodological approach that accounts for interdependent patient assignments in hospital resource allocation. Second, it demonstrates how

combining causal forests with a minimax regret policy improves NCU utilization while enhancing patient outcomes. The findings offer a practical strategy for reducing mortality and optimizing resource efficiency without increasing hospital capacity.

Data

To evaluate the effect of NCU placement on in-hospital mortality, I use administrative hospital data from the Swiss Federal Statistical Office's Medical Statistics of Hospitals. This dataset provides detailed information on patients' socio-demographic characteristics, outcomes, treatments coded according to the Swiss surgical classification (CHOP), and diagnoses coded using ICD-10-GM. Structured data collection began before the introduction of the Swiss Diagnosis-Related Groups (SwissDRG) reimbursement system in 2012 and covers all inpatient hospital stays in Switzerland. The dataset spans from 2012 to 2020, encompassing 278 hospitals and approximately 13 million cases.

We apply several sample restrictions. First, I focus on the five university hospitals. Second, to ensure a well-defined choice of NCU placement, I limit the sample to patients admitted either to the general internal medicine NCU or the surgical NCU, thereby reducing heterogeneity in admission options. Third, for clinical comparability, I restrict the sample to patients whose main diagnosis falls within specific ICD-10 groups: I2, I3, I6, I7, C1, C2, C3, C4, and C7. This results in a final analysis sample of *74,355* cases. While causal effects are estimated on the entire dataset, the policy is constructed for the years 2017 to 2019 for computational reasons.

Dependent Variable

The primary outcome of interest is extitin-hospital mortality. Our dependent variable is a binary indicator equal to one if the patient survived the hospital stay and zero otherwise. This study aims to examine how assignment to different NCUs affects this critical outcome and to derive insights for optimizing patient assignment policies (see Table **??** for variable definitions).

Independent Variables

The main independent variable of interest is the type of NCU to which the patient is admitted. This is recorded via the primary cost center assigned to each hospital stay, corresponding to the responsible NCU. I define an indicator variable $W = 1$ if the patient is admitted to the internal medicine NCU (treatment group) and $W = 0$ if admitted to the surgical NCU (control group).

Additionally, I control for a comprehensive set of patient and stay characteristics, including socio-demographic variables (e.g., age, sex), primary and secondary diagnoses (ICD-10), procedures performed (CHOP codes), as well as fixed effects for day-of-week, month, and year to account for temporal trends. Hospital fixed effects are included to control for unobserved, time-invariant hospital characteristics. At the time of admission, I assume that only the main diagnosis, diagnoses listed in the Elixhauser comorbidity score, and main treatments are known[1].

[1] Other diagnoses and treatments may be a result of complications or conditions that arise during the hospital stay

Instrumental Variable

To address potential selection bias in NCU assignment, I use daily utilization as an instrumental variable. Higher internal medicine utilization is expected to decrease the likelihood of admission to internal medicine and increase the likelihood of redirection to the surgical NCU, thereby serving as a valid instrument. To prevent mechanical correlations between the instrument and patient outcomes, the index patient is excluded from the calculation of utilization.

Ideally, I would measure the actual number of available beds or nurse workload in the ICU. However, such data are either unavailable or poorly recorded in our dataset. Consequently, our utilization measure is based solely on the daily count of patients, serving as a proxy for NCU utilization and capacity constraints at the time of admission.

The instrumental variable Z captures exogenous variation in NCU busyness. Specifically, I define the binary variables $Z_{Int.Med}$ and $Z_{Surg.}$ based on the median busyness levels of internal medicine and surgical NCUs, respectively, within each hospital and year. Given inter-hospital variation and temporal trends (as observed in Fig. 3), I calculate the median busyness of the internal medicine NCU for each group. $Z_{Int.Med}$ is set to 1 if the busyness exceeds the group median and 0 otherwise. Similarly, $Z_{Surg.}$ is set to 1 if the busyness is below the group median and 0 otherwise. This approach ensures that the instrument reflects relative busyness levels within each hospital and year, providing a robust proxy for capacity constraints independent of individual patient characteristics and outcomes.

D	Busyness Int.Med	Busyness Surg.	Z
1	$B_{IM}(h,t) > P_{50}(h,t)$	$B_S(h,t) > P_{50}(h,t)$	1
1	$B_{IM}(h,t) > P_{50}(h,t)$	$B_S(h,t) < P_{50}(h,t)$	1
1	$B_{IM}(h,t) < P_{50}(h,t)$	$B_S(h,t) > P_{50}(h,t)$	0
1	$B_{IM}(h,t) < P_{50}(h,t)$	$B_S(h,t) < P_{50}(h,t)$	0
0	$B_{IM}(h,t) > P_{50}(h,t)$	$B_S(h,t) > P_{50}(h,t)$	0
0	$B_{IM}(h,t) > P_{50}(h,t)$	$B_S(h,t) < P_{50}(h,t)$	1
0	$B_{IM}(h,t) < P_{50}(h,t)$	$B_S(h,t) > P_{50}(h,t)$	0
0	$B_{IM}(h,t) < P_{50}(h,t)$	$B_S(h,t) < P_{50}(h,t)$	1

Table 1: IV Operationalization

Methods

In the first step, I employ an instrumental variable causal forest to estimate the heterogeneous treatment effects of NCU admission on patient outcomes. This method accounts for potential confounding factors by exploiting exogenous variation in daily utilization, thereby providing reliable patient-specific policy scores. These scores quantify the expected benefit for each patient when admitted to a less congested yet appropriately specialized NCU.

The heterogeneous causal effects of admission to a particular NCU on patient outcomes are estimated using an instrumental variable causal forest, implemented via the `instrumental_forest` function from the `grf` package. This method flexibly models non-linear relationships and treatment effect heterogeneity by leveraging exogenous variation in treatment induced by fluctuations in NCU busyness.

In the second step, I develop a policy framework that optimizes patient placement in a hospital NCU using a minimax regret decision rule. This criterion seeks to minimize the worst-case regret across all possible patient assignments by explicitly considering the impact of one patient's placement on NCU utilization and, consequently, on the outcomes of other patients.

Setup and Identification for Instrumental Variables

To estimate the heterogeneous treatment effects of NCU admission on patient outcomes while addressing potential confounding, I employ an instrumental variable (IV) causal forest approach. The IV method exploits exogenous variation in NCU utilization, enabling robust causal inference. Let Y_i denote the outcome for patient i, D_i the treatment indicator (admission to a specific NCU), and Z_i an instrumental variable influencing D_i but not directly affecting Y_i. The standard potential outcomes framework defines the individual treatment effect as (Rubin, 1974):

$$\tau_i = Y_i(1) - Y_i(0),$$

where $Y_i(1)$ and $Y_i(0)$ are potential outcomes under treatment and control, respectively. However, in an instrumental variable setting, I estimate the local average treatment effect (LATE) rather than the average treatment effect (ATE), conditional on compliance with the instrument.
To ensure the validity of Z_i as an instrumental variable, the following identifying assumptions must hold (Angrist and Pischke, 2009):

1. **Instrument Relevance:** The instrument Z_i must be correlated with the treatment D_i, meaning that the busyness levels captured by Z should influence the likelihood of a patient being assigned to a specific NCU. Mathematically, this is expressed as:

   $$\mathbb{E}[D_i \mid Z_i = 1] \neq \mathbb{E}[D_i \mid Z_i = 0].$$

   In our context, higher busyness in the internal medicine NCU (or lower busyness in the surgical NCU) affects the probability of a patient being assigned to one unit over the other. While the instrument strength is decent but not exceptional, controlling for patient and hospital characteristics should help ensure the identification of causal effects.

2. **Instrument Exogeneity:** Conditional on the observed covariates X_i, the instrument Z_i must be independent of the potential outcomes $Y_i(1)$ and $Y_i(0)$. This ensures that Z does not directly affect the outcome except through its influence on the treatment D_i. Formally:

   $$Y_i(1), Y_i(0) \perp Z_i \mid X_i.$$

   This assumption is plausible because Z is constructed based on exogenous variation in NCU busyness, which is unrelated to individual patient characteristics. The use of busyness levels from the NCU to which the patient was not assigned further reinforces exogeneity.

3. **Exclusion Restriction:** The instrument Z_i affects the outcome Y_i only through its effect on the treatment D_i. This implies that any impact of Z on patient outcomes is mediated entirely by assignment to a specific NCU. Formally:

   $$\mathbb{E}[Y_i \mid Z_i, X_i] = \mathbb{E}[Y_i \mid D_i, X_i].$$

   In our setting, this means that busyness levels captured by Z influence patient outcomes only through their effect on NCU assignment. Similar to the instrument exogeneity assumption, this is ensured by using the busyness of the NCU to which the patient was not admitted.

4. **Monotonicity:** The effect of the instrument on treatment assignment must be monotonic. This means that for all patients, an increase in Z (e.g., higher busyness in internal medicine) either increases or does not change the likelihood of being assigned to the surgical NCU. Formally:
$$D_i(Z_i = 1) \geq D_i(Z_i = 0) \quad \forall i.$$
This assumption ensures that there are no defiers (patients who would move in the opposite direction of the instrument's effect). As in every observational study, some defiers exist, as shown in Fig. 5, where negative compliance score values appear. However, their presence is not substantial enough to introduce computational issues.

These assumptions collectively ensure that Z is a valid instrument, allowing us to estimate the causal effect of NCU assignment on patient outcomes. By leveraging exogenous variation in NCU busyness, Z provides a robust basis for addressing potential confounding and selection bias in the analysis.

ATE, CATE, and IATE Estimation

The estimation of the ATE and the CATE is based on the augmented inverse-propensity weighted scores by (Robins et al., 1994) and the causal forest developed by (Wager and Athey, 2018b; Athey et al., 2019; Chernozhukov et al., 2018), which is based on (Breiman, 2001).

$$\hat{\tau} = \frac{1}{n} \sum_{i=1}^{n} \hat{\Gamma}_i + \hat{\Psi}_i \tag{1}$$

$$\hat{\Gamma}_i = \hat{\mu}_{(1)}(X_i) - \hat{\mu}_{(0)}(X_i) + \frac{d_i}{\hat{e}(X_i)} \big(Y_i - \hat{\mu}_{(1)}(X_i)\big) - \frac{1 - d_i}{1 - \hat{e}(X_i)} \big(Y_i - \hat{\mu}_{(0)}(X_i)\big) \tag{2}$$

$$\hat{\Psi}_i = \frac{Z_i - \hat{Z}_i(X_i)}{\hat{Z}_i(X_i)(1 - \hat{Z}_i(X_i))} (Y_i - \hat{\mu}(X_i) * \hat{\Gamma}_i(d_i - \hat{e}(X_i))) \tag{3}$$

$\hat{\Gamma}_i$ consists of two components (Wager and Athey, 2018b):

- The estimated conditional non-parametric expected values of the outcome for the treated and non-treated group, $\hat{\mu}_{(1)}(X_i)$ and $\hat{\mu}_{(0)}(X_i)$.
- The estimated non-parametric propensity scores $\hat{e}(X_i)$, i.e. the probability that a patient will receive the treatment depending on the confounders.

$\hat{\Psi}_i$ debiases $\hat{\Gamma}_i$ using the instrumental variable and consists of two components:

- The estimated conditional non-parametric expected values of the outcome, $\hat{\mu}(X_i)$.
- The estimated conditional non-parametric expected probabilities of the instrument, $\hat{Z}_i(X_i)$.

The instrumental variable causal forest partitions the data into subgroups with similar covariate profiles and estimates heterogeneous treatment effects flexibly. This approach accommodates non-linear relationships and complex interactions between covariates, offering a robust alternative to traditional IV estimation methods. $\hat{\Gamma}_i$ is the estimated double robust score for patient i, which captures the expected difference in outcomes between the internal medicine and surgical NCUs in the selection of observable setting. $\hat{\Psi}$ corrects the estimand using the instrument. By leveraging exogenous variation in NCU occupancy, our methodology provides patient-specific treatment effect estimates, aiding in optimal patient assignment decisions while ensuring unbiased causal inference. All components are estimated through the causal forest, for which I use the grf package in R by Athey et al. (2019).

Policy Framework

In our setting the average treatment effect does not provide any usefull interpretation because the treatment effect is not constant across the patients. The treatment effect is heterogeneous and depends on the patient characteristics, the NCU busyness and the NCU specialization. Additionally, not all patients can be placed into one NCU, because of bed capacity. Therefore, a policy that optimises NCU placement is needed. Based on the rapidly expanding literature on optimal treatment assignment started by Manski (2004) and developed further by Hirano and Porter (2009) and Stoye (2009), Kitagawa and Tetenov (2018) developed a non-parametric solution for optimal policy assignment with known propensities. I extend the Mini-Max regret into a case with multiple decisons affecting multiple patients at once. The proposed policy aims to optimize the daily assignment of P admitted patients to two different NCUs. The primary trade-off is between assigning patients to their optimal specialized NCU and mitigating the negative effects of congestion, which increases mortality risk. The policy seeks to minimize the worst-case regret across all possible assignment configurations, ensuring that no alternative placement would have resulted in significantly better outcomes for any patient group.

Definition of Regret

Regret R in decision theory measures the loss incurred when a suboptimal choice is made compared to the best possible alternative (Savage, 1951). In our setting, the regret for assigning patient i to NCU $d_i \in \{0,1\}$ is defined as:

$$R_i(d_i) = \max_{d_i' \in \{0,1\}} \tilde{Y}_i(d') - \tilde{Y}_i(d) \tag{4}$$

where, the expected outcome under a placement $\tilde{Y}_i(d)$ is needed instead of the treatment effect:

$$\tilde{Y}_i(d) = \hat{\mu}_{(d)}(X_i) + \frac{1}{\hat{e}_d(X_i)} \left(Y_i - \hat{\mu}_{(d)}(X_i)\right) \mathbb{1}_{D=d} + \Psi_i(d_i) \tag{5}$$

where:

- $Y_i(d_i)$ represents the potential outcome of patient i when assigned to NCU d_i, where $d_i = 1$ indicates assignment to the Internal Medicince NCU and $d_i = 0$ indicates assignment to the surgical NCU.
- $\max_{d_i' \in \{0,1\}} Y_i(d_i')$ represents the best possible outcome patient i could have achieved under an alternative assignment.

This regret formulation captures the difference in patient outcomes between the chosen and the optimal assignment.

Since P patients are assigned simultaneously each day, the regret function must account for all patient assignments. Let $\boldsymbol{\pi} = (\pi_1, \pi_2, ..., \pi_P)$ denote a specific assignment configuration, then the total regret for a given configuration is:

$$\mathrm{R}(\boldsymbol{\pi}) = \sum_{i=1}^{P} \left[\max_{d_i' \in \{0,1\}} \tilde{Y}_i(\pi') - \tilde{Y}_i(\pi) \right] \tag{6}$$

where:

- π is the vector of patient assignments across the P patients admitted on a given day.
- The summation aggregates regret across all patients to assess the overall regret of the daily assignment decision.

Minimax Regret Decision Rule

The minimax regret policy seeks to minimize the maximum possible regret across all possible assignments, ensuring that the worst-case regret is as small as possible. Formally, the optimal assignment configuration $\boldsymbol{\pi}^*$ is:

$$\boldsymbol{\pi}^* = \arg\min_{\mathbf{d} \in \boldsymbol{\Pi}} \max_{\pi' \in \boldsymbol{\Pi}} \left[\sum_{i=1}^{P} \tilde{Y}_i(\pi') - \tilde{Y}_i(\pi) \right], \tag{7}$$

where:

- $\diamond$ represents the set of all feasible assignment configurations for the P patients.
- π^* is the assignment configuration that minimizes the worst-case regret across all possible assignments.
- $\max_{\pi' \in \boldsymbol{\Pi}} \left[\sum_{i=1}^{P} \tilde{Y}_i(d') - \tilde{Y}_i(d) \right]$ captures the worst-case regret for any possible alternative assignment.

Welfare Effects of the Policy

The welfare generated by the policy is evaluated as the expected sum of patient outcomes under the chosen assignment:

$$W(\mathbf{d}) = \sum_{i=1}^{P} \tilde{Y}_i(\pi) \tag{8}$$

To incorporate treatment effect heterogeneity, I express the welfare as:

$$W(\boldsymbol{\pi}) = \sum_{i=1}^{P} \left\{ \tilde{Y}_{ii}(0) + (\hat{\Gamma}_i + \hat{\Psi}_i)\pi_i \right\} \tag{9}$$

where:

- $Y_i(0)$ is the potential outcome of patient i if assigned to the Internal Medicine NCU ($\pi_i = 0$).

- $\tau(X_i)$ is the conditional average treatment effect (CATE) for patient i, which measures the difference in expected outcomes between the specialized and general NCUs.
- The second summation quantifies the total treatment effect across all assigned patients.

This equation ensures that the policy not only selects the best individual assignments but also considers system-wide effects on congestion. The welfare can be interpreted as the expected number of patients that survive the hospital stay. $1 - W$ would be the number of patients that died during their stay.

Implementation

The policy operates in the following steps:

1. **Daily Patient Assignment:** Each day, P patients arrive and must be assigned to one of the two NCUs.
2. **Outcome Prediction:** Use the instrumental variable causal forest to predict CATEs and IATEs for each patient under both possible assignments.
3. **Regret Calculation:** Compute regret for each possible assignment by comparing predicted outcomes against the best feasible assignment.
4. **Decision Rule:** The optimal assignment is determined by finding the configuration $\boldsymbol{\pi}^*$ that minimizes the maximum regret across all assignments.
5. **Assignment Implementation:** The optimal allocation is implemented by updating patient records with the best NCU placement.

This policy ensures that patient assignment decisions are robust to worst-case scenarios, balancing individual treatment effects with system-wide constraints.

2 Results

In the following section, I first show descriptive statistics and the ATE of the NCU placement comparing internal medicine and the surgical NCU. The I show the CATE by busyness in the internal medicine and the surgical NCU. In the end I show the new utilization over time and the expected saved mortality.

2.1 Descriptives

Metric	Overall	Int.Med	Surg.
N	74355	44716	29639
Survival (Y)	0.96	0.96	0.97
Int.Med. (W)	0.60	1.00	0.00
$B_{P_{50}}$ (Z)	0.47	0.48	0.46

Table 2: Summary statistics for N, Y, W, and Z.

The Table (2) provides an overview of the dataset, summarizing key patient characteristics, and NCU placements. Table 15 and Table 13 in the Appendix summarizes the main diagnosis of the patients and other all control variables. 60% of the patient in my sample are placed in the internal medicine NCU and the other 40% in the surgical NCU. 96% survive their stay with a overall mortality rate of 4% with minor variation between the two NCUs.

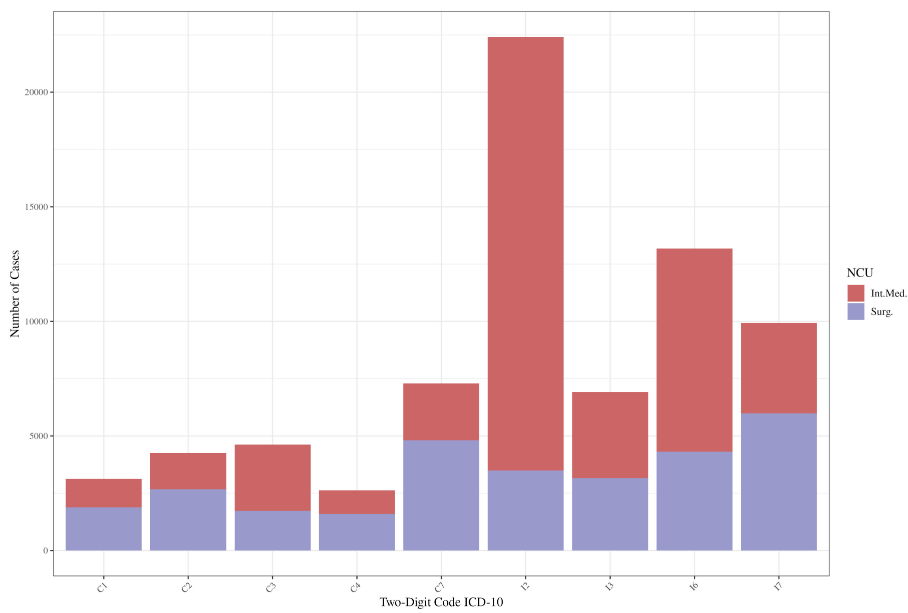

Figure 1: Stacked Bar Plot of Cases by Two-Digit ICD-10 Codes and NCU Placement

The stacked bar plot (Figure 1) provides a visual representation of the distribution of cases across two-digit ICD-10 codes and their corresponding NCU placements. The x-axis displays the two-digit ICD-10 codes, which categorize patients based on their primary diagnoses, while the y-axis represents the total number of cases for each diagnosis. Each bar is divided into two segments, with red indicating cases assigned to internal medicine NCUs and blue representing cases assigned to surgical NCUs.

The plot highlights significant variation in the distribution of cases across ICD-10 codes, reflecting the influence of clinical specialization and capacity constraints on patient placement decisions. For instance, ICD-10 code I2, which corresponds to ischemic heart diseases, accounts for the largest number of cases, with approximately 60% of these cases assigned to internal medicine NCUs and 40% to surgical NCUs. In contrast, ICD-10 code C3, representing malignant neoplasms, shows a different pattern, with 70% of cases assigned to surgical NCUs and 30% to internal medicine NCUs. Other codes, such as C1 and C2, exhibit a more balanced distribution between the two types of NCUs, indicating that placement decisions for these diagnoses are less influenced by specialization. The plot also reveals that certain diagnoses, such as I2 and I6, are predominantly treated in internal medicine NCUs, while others, such as C3 and C4, are more frequently managed in surgical NCUs. This variation underscores the importance of aligning patient needs with the specialized capabilities of NCUs to optimize outcomes. Furthermore, the stacked bar plot provides insights into the overall workload distribution between internal medicine and surgical NCUs, highlighting the need for effective patient placement policies to balance resource utilization and clinical outcomes.

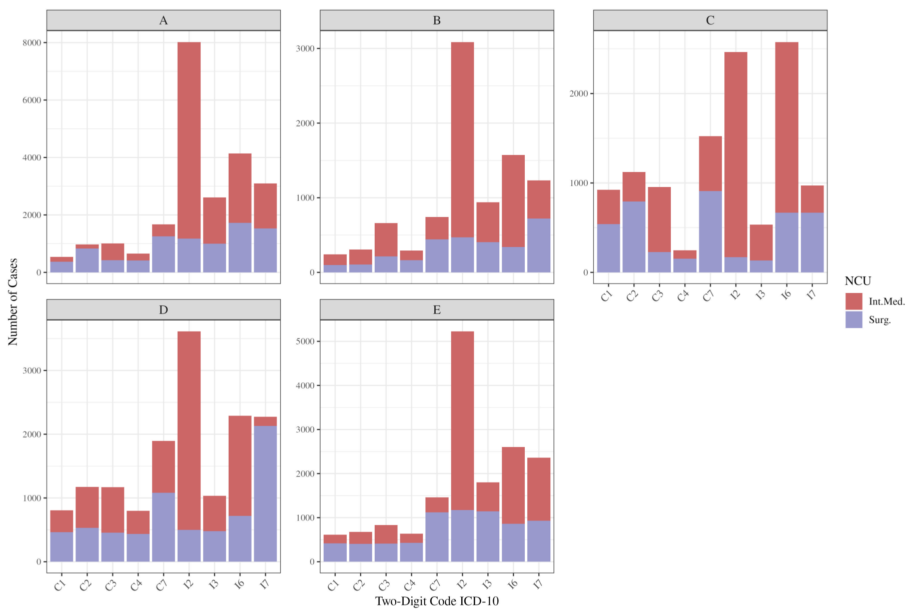

Figure 2: Cases by Two-Digit ICD-10 Codes, NCU Placement and Hospital

The panel stacked bar plot (Figure 2) provides a more granular view of the distribution of cases across two-digit ICD-10 codes and NCU placements, segmented by Hostpital. While the stacked bar plot (Figure 1) offers an aggregated perspective of the overall distribution, this panel plot breaks down the data by canton, allowing for a detailed comparison of regional variations in patient placement decisions. Each panel corresponds to a specific Hospital, with the x-axis representing the two-digit ICD-10 codes and the y-axis showing the total number of cases for each diagnosis. Similar to the stacked bar plot, the bars are divided into two segments, with red indicating cases

assigned to internal medicine NCUs and blue representing cases assigned to surgical NCUs.
The panel plot reveals notable regional differences in the distribution of cases and NCU placements. For example, in Hosptital A, ICD-10 code I2 (ischemic heart diseases) accounts for a significant proportion of cases, with the majority assigned to internal medicine NCUs. In contrast, hospital C shows a more balanced distribution of cases between internal medicine and surgical NCUs for the same diagnosis. Similarly, while ICD-10 code C3 (malignant neoplasms) is predominantly managed in surgical NCUs across most cantons, the proportion of cases assigned to internal medicine NCUs varies, reflecting differences in regional capacity and specialization.
By comparing Figures 1 and 2, it becomes evident that the overall trends observed in the aggregated data mask important regional variations. For instance, while the stacked bar plot indicates that ICD-10 code I6 (heart failure) is predominantly treated in internal medicine NCUs, the panel plot shows that this pattern is not consistent across all cantons. Such as Hospital D, exhibit a higher proportion of cases assigned to surgical NCUs for this diagnosis, likely due to differences in local resource availability or clinical practices.

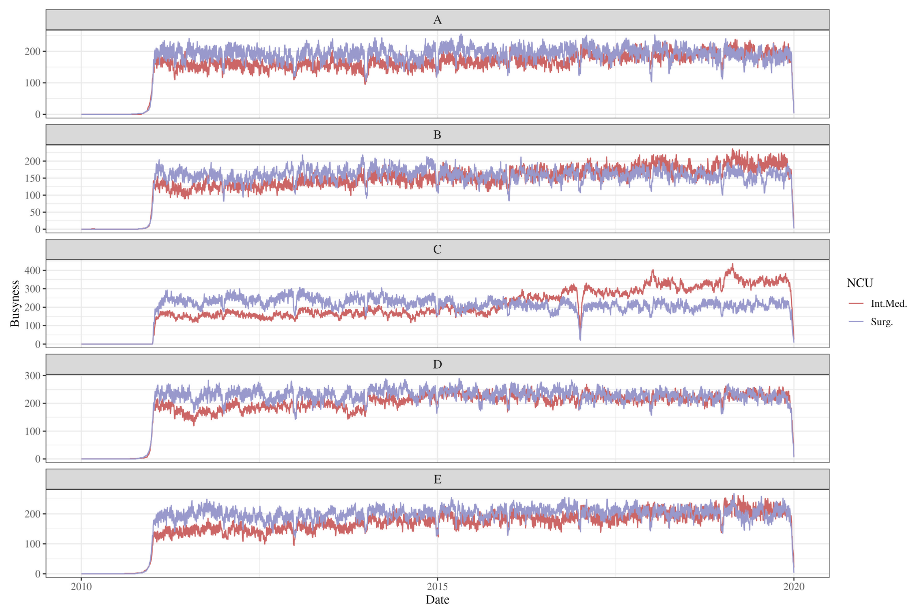

Figure 3: Time Series of Busyness by Hospital

The time series plot (Figure 3) illustrates the variation in busyness levels over time for each Hospital. The x-axis represents the timeline spanning from 2010 to 2020, while the y-axis shows the busyness levels, measured as the number of patients in the NCU. Each panel corresponds to a specific Hospital, providing a detailed view of the temporal trends in NCU utilization. The plot captures fluctuations in busyness levels across the years, with distinct patterns observed for each canton. There is a general trend that more patients with the selected diagnosis groups are beeing placed in the internal medicine NCU. This trend can be simpliy an increase in the prevalence of patients with a specific Diagnosis that are predominately placed in the internal medicine NCU, and do not necessarily mean that there operational changes in the hospitals.

2.2 IV Descriptives

Model	First Stage R2	First Stage F-statistic
Linear Model	0.38	99.5
Causal Forest	0.39	

Table 3: Comparison of First Stage Metrics for Linear Model and Instrumental Forest

Table **??** presents a comparison of first stage metrics for two models: the Linear Model (LM) and the Causal Forest (CF). The table includes four key metrics: the First Stage R-squared, the Partial R-squared, and the First Stage F-statistic. For the Linear Model (LM), the First Stage R-squared is reported as 0.38, and the First Stage F-statistic is 99.5. For the Causal Forest , the First Stage R-squared is 0.39. [2] While the instrument strength is decent but not exceptional, controlling for patient and hospital characteristics should help ensure the identification of causal effects.

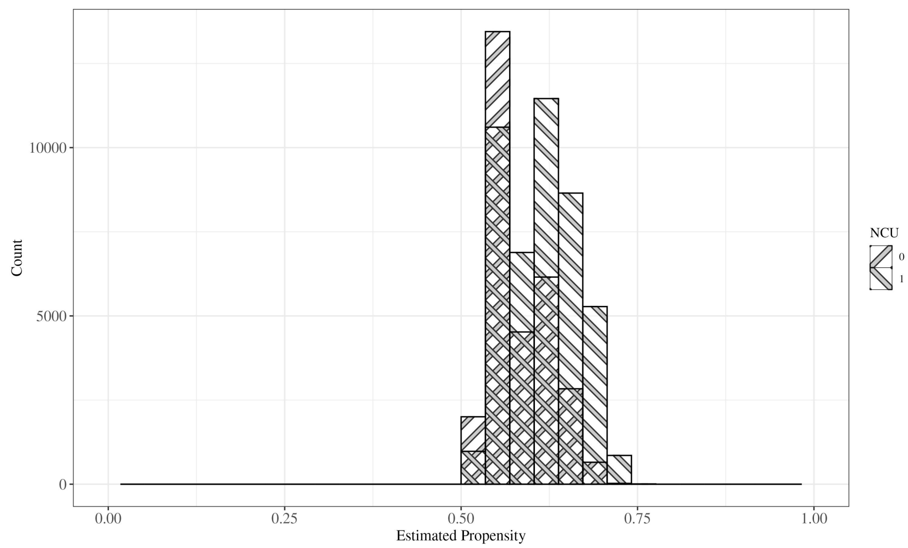

Figure 4: Common Support Plot

The common support plot (Figure 4) illustrates the distribution of estimated propensity scores for patients assigned to different NCUs. The x-axis represents the estimated propensity scores, ranging from 0 to 1, while the y-axis indicates the count of patients. The plot is divided into two groups based on NCU assignment, with one group corresponding to patients assigned to internal medicine NCUs (indicated by NCU = 1) and the other group corresponding to patients assigned to surgical NCUs (indicated by NCU = 0). The histogram uses overlapping patterns to display the distribution of propensity scores for both groups, providing a visual representation of the common support region where the propensity scores overlap between the two groups. It shows that there is no strong selection between the NCU, there I can assume that common support is fulfilled.

[2] The First Stage F-statistic is not reported for the Causal Forest model.

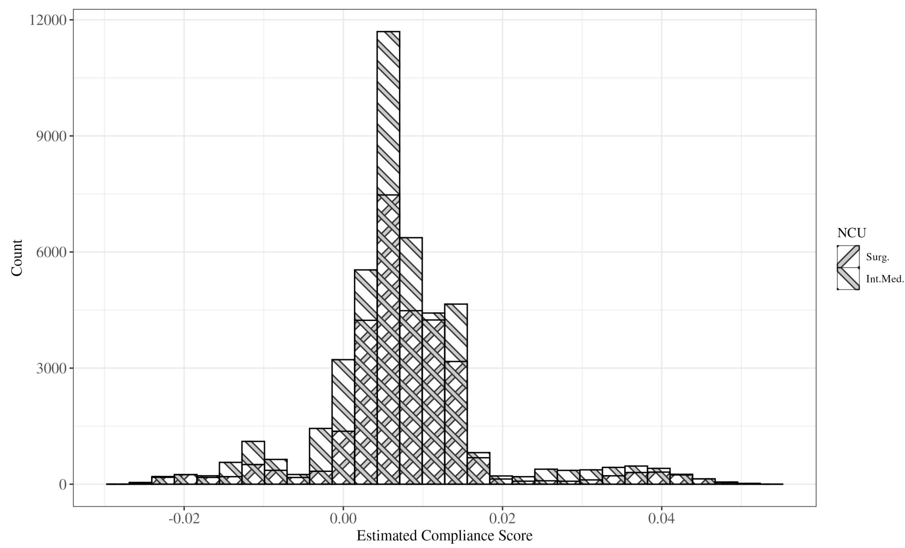

Figure 5: Compliance Score Descriptive Plot

The compliance score descriptive plot (Figure 5) illustrates the distribution of estimated compliance scores for patients assigned to different NCUs. The x-axis represents the estimated compliance scores. The plot is divided into two groups based on NCU assignment, with one group corresponding to patients assigned to internal medicine NCUs and the other group corresponding to patients assigned to surgical NCUs. The compliance scores are displayed as a histogram with overlapping bars for the two groups, providing a visual representation of the distribution of compliance scores across the patient population. Most observations have a positive values in their compliance scores, with some negative values. These are most likely simply some statistical artifacts.

2.3 (Conditional) Average Treatment Effects

	Model	Average.Treatment.Effect	Standard.Error
estimate	Causal Forest	-0.01	0.00
W_hat	Linear Model	-0.02	0.00

Table 4: Average Treatment Effect

Table 5: Average Treatment Effect (ATE) Table

Table 5 presents the average treatment effect (ATE) estimates for two models: the Causal Forest and the Linear Model. The table includes the estimated ATE and the corresponding standard error for each model. For the Causal Forest model, the estimated ATE is -0.01 with a standard error of 0.00. For the Linear Model, the estimated ATE is -0.02 with a standard error of 0.00. Both effects are significant with $\alpha = 0.05$ and show the same direction. With an overall mortality of 4% The effect can be interpreted as increase in the mortality rate of 25% being placed in the internal medicine compared to the surgical NCU. As we assume that there is large heterogeneity the effect has no practical relevance. Additionally, this effect is most likely driven by patients that should are and should be treated in the surgical NCU.

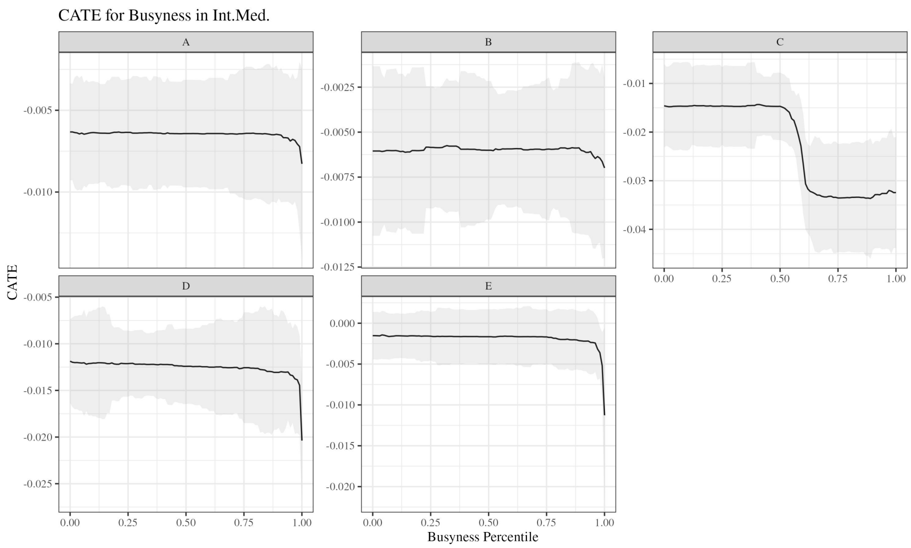

Figure 6: Conditional Average Treatment Effects (CATE) for Busyness 100

The conditional average treatment effects (CATE) plot for busyness at 100% (Figure 6) illustrates the predicted treatment effects across different cumulative distribution function (CDF) values for five groups of patients. The x-axis represents the CDF values, ranging from 0 to 1, while the y-axis shows the predicted treatment effects. Each panel corresponds to a specific group, with shaded regions indicating the confidence intervals around the predicted effects. The plot provides a detailed view of how treatment effects vary with an increasing busyness, for most of the distribution there are minor changes in the effect, however in the 10 percentiles with the highest busyness the effect increases.

2.4 Policy Implications

	Policy = 0	Policy = 1
D = 0	5939	4994
D = 1	7836	9626

Table 6: Confusion Matrix between W and Policy Hospital

Table 6 presents the confusion matrix comparing the observed NCU placements (W) with the policy-assigned placements (*Policy Hospital*). The rows represent the observed placements, while the columns represent the policy-assigned placements. The diagonal elements of the matrix indicate cases where the observed and policy-assigned placements align, while the off-diagonal elements capture discrepancies between the two placement strategies. Specifically, the table shows that 5,939 patients were observed in $W = 0$ (e.g., surgical NCU) and were also assigned to *Policy = 0* by the policy, while 9,626 patients were observed in $W = 1$ (e.g., internal medicine NCU) and were assigned to *Policy = 1*. These values represent the cases where the policy agrees with the observed placements. In contrast, 4,994 patients were observed in $W = 0$ but were assigned to *Policy = 1*, and 7,836 patients were observed in $W = 1$ but were assigned to *Policy = 0*. These discrepancies highlight the potential reallocation of patients under the policy framework, which aims to optimize outcomes by balancing specialization and utilization.

NCU	Max_Δ	Min_Δ	μ_Δ	sd_Δ
Int.Med.	20.00	-23.00	-0.52	2.45
Surg.	23.00	-20.00	0.52	2.45

Table 7: Summary of Differences in Busyness Between Policy and Original

Table 7 provides an aggregated summary of the differences in busyness levels between the policy-assigned and observed placements for the two types of NCUs: internal medicine (*Int.Med.*) and surgical (*Surg.*). The table reports the maximum, minimum, mean, and standard deviation of the differences in busyness levels. For internal medicine NCUs, the maximum difference is 20, while the minimum difference is -23, with a mean difference of -0.52 and a standard deviation of 2.45. For surgical NCUs, the maximum difference is 23, the minimum difference is -20, with a mean difference of 0.52 and a standard deviation of 2.45. These results indicate that the policy framework introduces small but balanced changes in busyness levels across the two NCU types, with a slight reduction in internal medicine busyness and a corresponding increase in surgical busyness. Table 8 extends this analysis by grouping the differences in busyness levels by hospital (*UH Canton*). The table provides a detailed breakdown of the maximum, minimum, mean, and standard deviation of the differences for each hospital and NCU type. For example, in hospital A, the mean difference in busyness for internal medicine NCUs is 0.00, while for surgical NCUs, it is -0.00, with a standard deviation of 1.31 for both. In contrast, hospital E exhibits the largest variation, with a mean difference of -2.00 for internal medicine NCUs and 2.00 for surgical NCUs, and a standard deviation of 4.71 for both. These results highlight the heterogeneity in the impact of the policy across hospitals, reflecting differences in baseline utilization levels and patient placement patterns. Together, these tables provide insights into how the policy framework redistributes busyness levels across NCUs and hospitals, aiming to balance utilization while minimizing disruptions to existing workflows.
Table 9 provides a summary of busyness levels for internal medicine (*Int.Med.*) and surgical (*Surg.*) NCUs under both the observed (Original) and policy-assigned (Policy) placements. The table reports the maximum, minimum, and mean busyness levels for each NCU type. For internal medicine NCUs, the maximum busyness under the observed placements is 436, while under the policy placements it is slightly reduced to 434. The mean busyness for internal medicine NCUs decreases marginally from 222.11 to 221.59 under the policy. Conversely, for surgical NCUs, the maximum busyness increases from 270 to 277 under the policy, with the mean busyness rising from 196.87 to 197.39. These results indicate that the policy framework redistributes busyness levels, slightly reducing internal medicine busyness while increasing surgical busyness.
Table 10 extends this analysis by grouping the busyness levels by hospital (*UH Canton*). The

Hospital	NCU	Max_Δ	Min_Δ	μ_Δ	sd_Δ
A	Int.Med.	6.00	-6.00	0.00	1.31
A	Surg.	6.00	-6.00	-0.00	1.31
B	Int.Med.	6.00	-6.00	-0.17	1.00
B	Surg.	6.00	-6.00	0.17	1.00
C	Int.Med.	3.00	-6.00	-0.35	0.98
C	Surg.	6.00	-3.00	0.35	0.98
D	Int.Med.	6.00	-6.00	-0.07	1.14
D	Surg.	6.00	-6.00	0.07	1.14
E	Int.Med.	20.00	-23.00	-2.00	4.71
E	Surg.	23.00	-20.00	2.00	4.71

Table 8: Summary of Differences in Busyness Between Policy and Original, Grouped by by Hospital

NCU	Placement	Max_B	Min_B	μ_B
Int.Med.	Original	436.00	8.00	222.11
Int.Med.	Policy	434.00	8.00	221.59
Surg.	Original	270.00	7.00	196.87
Surg.	Policy	277.00	7.00	197.39

Table 9: Summary of Busyness by NCU with and without Policy

table provides a detailed breakdown of the maximum, minimum, and mean busyness levels for each hospital and NCU type under both observed and policy placements. For example, in hospital A, the mean busyness for internal medicine NCUs remains nearly unchanged, with values of 191.03 under observed placements and 191.04 under policy placements. For surgical NCUs in hospital A, the mean busyness remains constant at 193.34 under both observed and policy placements. In contrast, hospital E exhibits more pronounced changes, with the mean busyness for internal medicine NCUs decreasing from 201.82 to 199.82 under the policy, while the mean busyness for surgical NCUs increases from 203.91 to 205.92. These results highlight the heterogeneity in the impact of the policy across hospitals, reflecting differences in baseline utilization levels and patient placement patterns.

Together, these tables provide insights into how the policy framework redistributes busyness levels across NCUs and hospitals, aiming to balance utilization while minimizing disruptions to existing workflows.

	Hospital	NCU	Placement	Max_B	Min_B	μ_B
1	A	Int.Med.	Original	238.00	9.00	191.03
2	A	Int.Med.	Policy	237.00	9.00	191.04
3	A	Surg.	Original	252.00	12.00	193.34
4	A	Surg.	Policy	254.00	12.00	193.34
5	B	Int.Med.	Original	236.00	8.00	180.90
6	B	Int.Med.	Policy	233.00	8.00	180.73
7	B	Surg.	Original	201.00	7.00	158.58
8	B	Surg.	Policy	201.00	7.00	158.76
9	C	Int.Med.	Original	436.00	36.00	315.17
10	C	Int.Med.	Policy	434.00	36.00	314.82
11	C	Surg.	Original	266.00	15.00	205.05
12	C	Surg.	Policy	266.00	15.00	205.40
13	D	Int.Med.	Original	268.00	18.00	221.62
14	D	Int.Med.	Policy	264.00	18.00	221.54
15	D	Surg.	Original	270.00	10.00	223.47
16	D	Surg.	Policy	269.00	10.00	223.55
17	E	Int.Med.	Original	261.00	32.00	201.82
18	E	Int.Med.	Policy	270.00	32.00	199.82
19	E	Surg.	Original	268.00	7.00	203.91
20	E	Surg.	Policy	277.00	7.00	205.92

Table 10: Summary of Busyness by NCU and by Hospital with and without Policy

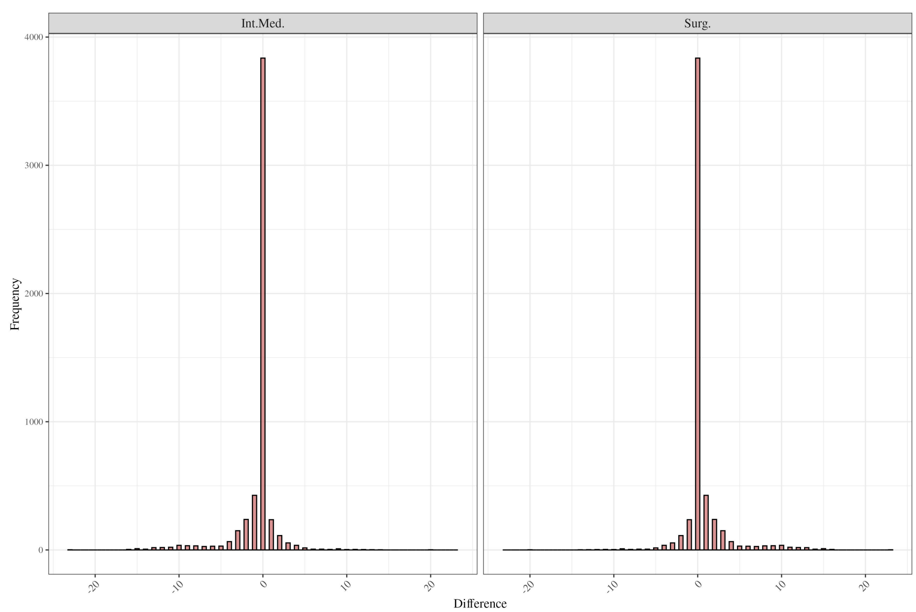

Figure 7: Busyness difference

Figure **??** illustrates the distribution of differences in busyness levels between the policy-assigned and observed placements for the two types of NCUs: internal medicine (*Int.Med.*) and surgical (*Surg.*). The x-axis represents the difference in busyness levels, while the y-axis shows the frequency of occurrences. The histogram is faceted by NCU type, providing a clear comparison of how the policy framework impacts busyness levels across the two NCU types. The results indicate that the differences are centered around zero, with a slight redistribution of busyness levels. For internal medicine NCUs, the policy slightly reduces busyness, while for surgical NCUs, the policy slightly increases busyness. This balanced redistribution reflects the policy's aim to optimize resource utilization while minimizing disruptions.
Figure **??** extends this analysis by grouping the differences in busyness levels by hospital (*UH Canton*). The histogram is faceted by both hospital and NCU type, providing a more granular view of the policy's impact across different hospitals. Each panel corresponds to a specific hospital and NCU type, allowing for a detailed comparison of regional variations in the distribution of busyness differences. The results highlight that while the overall trends are consistent across hospitals, there is notable heterogeneity in the magnitude of the differences. For example, some hospitals exhibit larger shifts in busyness levels under the policy, reflecting differences in baseline utilization and patient placement patterns.
Together, these figures provide valuable insights into how the policy framework redistributes busyness levels across NCUs and hospitals, aiming to balance utilization while minimizing disruptions to existing workflows.

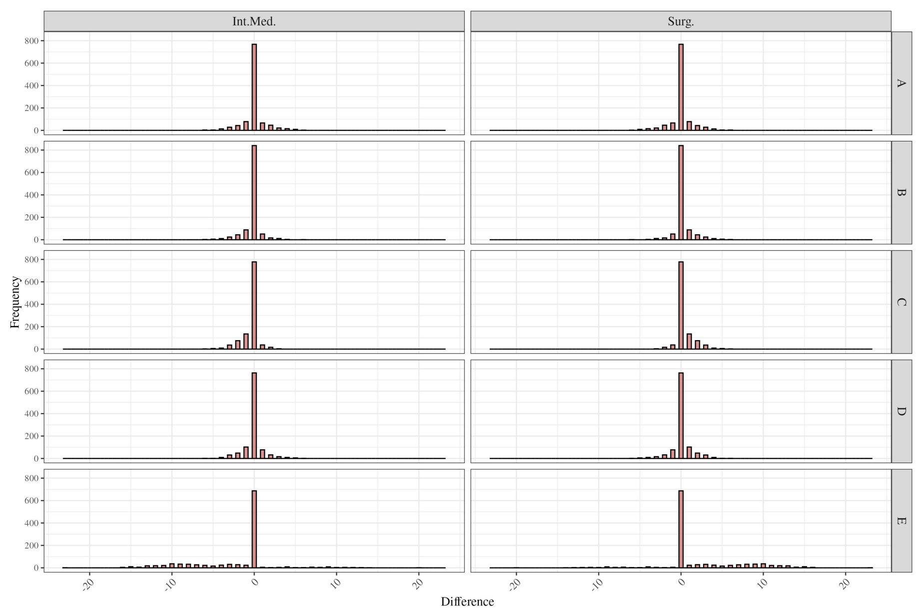

Figure 8: Busyness difference by hospital

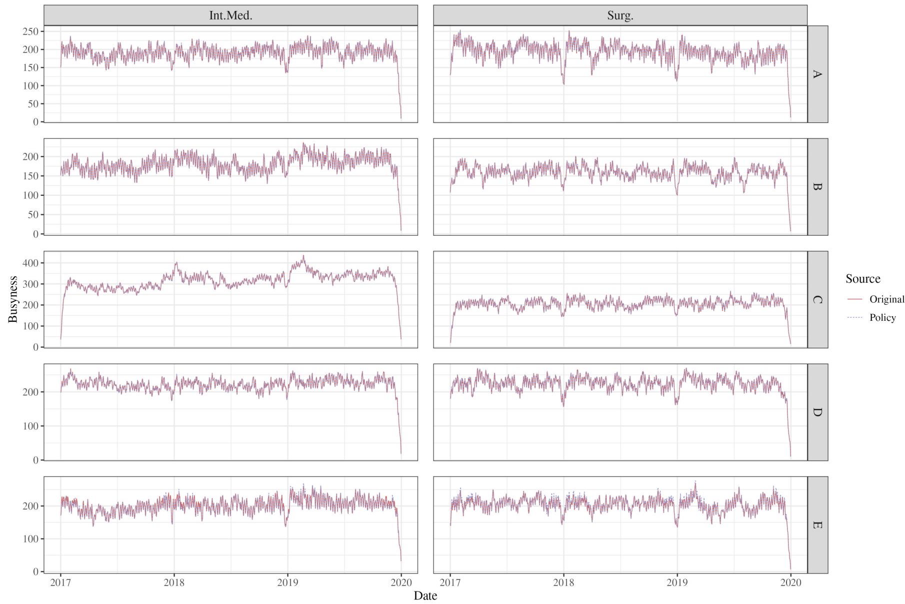

Figure 9: Add caption hear

Figure 9 illustrates the time-series of busyness levels for internal medicine (*Int.Med.*) and surgical (*Surg.*) NCUs across five hospitals (*UH Canton*) under both observed (Original) and policy-assigned (Policy) placements. The x-axis represents the timeline from 2017 to 2020, while the y-axis shows the busyness levels, measured as the number of patients in the NCU. Each panel corresponds to a specific hospital, with separate plots for internal medicine and surgical NCUs.

The figure highlights the temporal trends in busyness levels and compares the observed and policy scenarios. The red solid lines represent the observed busyness levels, while the blue dashed lines represent the policy-assigned busyness levels. Across all hospitals, the policy framework introduces minimal changes to the overall busyness trends, ensuring consistency with observed patterns. However, slight adjustments are visible, particularly in the redistribution of busyness between internal medicine and surgical NCUs, reflecting the policy's aim to balance utilization.

For example, in hospital A, the policy slightly reduces busyness in internal medicine NCUs while increasing it in surgical NCUs, maintaining overall stability. Similar patterns are observed in hospitals B and C, where the policy framework achieves a balanced redistribution of busyness levels. These results demonstrate the policy's ability to optimize resource utilization without causing significant disruptions to existing workflows.

Overall, the time-series plot provides valuable insights into the temporal dynamics of NCU utilization under the policy framework, highlighting its potential to improve resource allocation while maintaining operational stability.

Table 11 provides a summary of the welfare metrics under the observed and policy-assigned placements. The table includes total welfare, average welfare, relative change, and absolute change between the observed and policy scenarios. The total welfare under the observed placements is 1,190, while the policy reduces this slightly to 1,163.11, resulting in a relative change of 2% and an absolute change of -26.89. Similarly, the average welfare remains consistent at 0.04 under both

Metric	Observed	Policy	$\Delta_{\%}$	Δ_{abs}
Total Welfare	1190.00	1163.11	0.02	-26.89
Average Welfare	0.04	0.04	0.02	-0.00

Table 11: Comparison of Observed and Policy Welfare Metrics

observed and policy placements, with a relative change of 2% and an absolute change of -0.00. These results indicate that the policy framework introduces minor adjustments to welfare metrics while maintaining overall consistency.

Table 12 extends this analysis by grouping the welfare metrics by hospital (*UH Canton*). The table provides a detailed breakdown of total and average welfare, along with their relative and absolute changes, for each hospital. For example, in hospital A, the total welfare decreases slightly from 1,190.00 under observed placements to 1,185.64 under the policy, with a relative change of 0% and an absolute change of -4.36. Similarly, the average welfare remains consistent at 0.04, with a relative change of 1% and an absolute change of -0.00. In contrast, hospital C exhibits a more pronounced reduction in total welfare, decreasing from 1,190.00 to 1,170.90, with a relative change of 2% and an absolute change of -19.10. These results highlight the heterogeneity in the policy's impact across hospitals, reflecting differences in baseline utilization and patient placement patterns.

Together, these tables provide insights into the welfare effects of the policy framework, demonstrating its ability to optimize patient placement while maintaining overall welfare metrics across hospitals.

UH_canton	Observed_Total_Welfare	Policy_Total_Welfare	Relative_Total_Change	Absolute_Total_Change	Observed_A
BE	185.00	180.64	0.02	-4.36	
BS	98.00	94.95	0.03	-3.05	
GE	344.00	324.90	0.06	-19.10	
VD	212.00	208.46	0.02	-3.54	
ZH	351.00	354.15	-0.01	3.15	

Table 12: Comparison of Observed and Policy Welfare Metrics by UH Canton

3 Discussion

The findings of this study provide valuable insights into optimizing patient placement in NCUs by balancing the trade-off between utilization and specialization. Addressing the first research question—*How can a patient placement policy be designed to minimize mortality while balancing utilization and specialization?*—the results demonstrate that the proposed policy framework effectively redistributes patients across NCUs, achieving a balance between reducing congestion and maintaining specialization. The minimax regret decision rule ensures that patient assignments minimize the worst-case outcomes, thereby improving overall resource allocation and patient outcomes.

The time-series analysis (Figure 9) highlights that the policy framework introduces minimal disruptions to existing busyness trends while achieving a slight redistribution of busyness between internal medicine and surgical NCUs. This redistribution reflects the policy's ability to optimize resource utilization without compromising operational stability. The welfare analysis (Tables 11 and 12) further supports the effectiveness of the policy. While the total welfare under the policy is slightly reduced compared to the observed placements, the relative and absolute changes are minimal, indicating that the policy maintains overall welfare metrics. Importantly, the grouped analysis reveals heterogeneity in the policy's impact across hospitals, reflecting differences in baseline utilization and patient placement patterns. This underscores the importance of tailoring policies to local hospital contexts to maximize their effectiveness. Addressing the second research question—*What are the welfare effects of implementing such a policy?*—the results indicate that the policy framework achieves its objectives with minimal adverse effects on welfare. The confusion matrix (Table 6) highlights the extent of reallocation required under the policy, with a significant number of patients reassigned to optimize outcomes. Despite these reallocations, the policy maintains consistency with observed patterns, as evidenced by the small changes in welfare metrics. The analysis of busyness differences (Tables 7 and 8) provides additional insights into the policy's impact on resource utilization. The results show that the policy introduces small but balanced changes in busyness levels across NCUs, with a slight reduction in internal medicine busyness and a corresponding increase in surgical busyness. The grouped analysis reveals notable heterogeneity across hospitals, emphasizing the need for hospital-specific adjustments to the policy framework. These findings are consistent with the motivation outlined in the *Introduction*, which highlighted the critical role of patient placement in shaping clinical outcomes and managing hospital resources. By addressing the trade-off between utilization and specialization, the proposed policy framework offers a practical solution to optimize patient placement, reduce mortality, and improve resource allocation efficiency. The results demonstrate that data-driven policies can achieve these objectives without necessitating additional hospital capacity, making them particularly relevant in resource-constrained environments. Overall, this study contributes to the growing body of literature on hospital resource management by providing a robust methodological framework for optimizing patient placement. The integration of instrumental variable causal forests with a minimax regret decision rule represents a novel approach to addressing the interdependence of patient assignments, offering actionable insights for policymakers and hospital administrators.

Limitations

While our study provides valuable insights into optimizing patient placement in NCUs, several limitations must be acknowledged. (1) our dataset does not include information on actual staff availability in the NCUs. As a result, I cannot directly observe the true utilization levels, which might influence patient outcomes beyond the estimated effects. (2) our analysis is based on a subset of NCUs and patients. Due to overlapping admissions and constraints in the dataset, I do not assign all potential patients to NCUs, which may limit the generalizability of our findings. (3) I assume that the instrumental variable—whether a surgical or internal medicine NCU is fully occupied—affects treatment probability in the same way. However, differences in clinical practices and organizational structures between these unit types could lead to deviations from this assumption. (4) our study is limited to patients with specific ICD-10 diagnoses. While this selection allows for a focused analysis, it also restricts the applicability of our results to a broader patient population. Future research should investigate whether similar assignment policies yield comparable results for other diagnostic groups. (5) our findings are based on data from a limited number of hospitals. The results might not fully capture variations in hospital-specific policies, resources, or patient demographics. Extending the analysis to a wider range of hospitals could improve the external validity of our conclusions. Despite these limitations, our approach provides a meaningful step towards data-driven patient placement policies, with potential implications for improving clinical outcomes and optimizing hospital resource utilization.

Future Research

This study provides a foundation for optimizing patient placement in normal care units (NCUs), but several avenues for future research remain. First, the analysis could be extended to include a broader range of patients, NCUs, and hospitals. Expanding the dataset would improve the generalizability of the findings and provide insights into the applicability of the proposed policy framework across diverse healthcare settings. Second, alternative specifications for the instrumental variable Z could be explored. For example, incorporating additional measures of NCU busyness or other hospital-level factors might enhance the robustness of the causal inference and provide a more nuanced understanding of the relationship between utilization and patient outcomes. Third, future work could extend the evaluation of the policy framework by incorporating its effects on length of stay (LOS). Understanding how the placement policy influences LOS would provide valuable insights into its broader implications for hospital resource utilization and patient flow management. Finally, estimating the effects of the placement policy on LOS and hospital costs would allow for a more comprehensive evaluation of its economic impact. Quantifying these outcomes would enable policymakers and hospital administrators to assess the cost-effectiveness of the proposed framework and make informed decisions about its implementation. By addressing these directions, future research could build on the findings of this study to further improve patient outcomes and optimize hospital resource allocation.

4 Conclusion

References

Abir, M., Goldstick, J., Malsberger, R., Bauhoff, S., Setodji, C. M., and Wenger, N. (2020). The association between hospital occupancy and mortality among medicare patients. *The Joint Commission Journal on Quality and Patient Safety*, 46(9):506–515.

Alameda, C. and Suárez, C. (2009). Clinical outcomes in medical outliers admitted to hospital with heart failure. *European journal of internal medicine*, 20(8):764–767.

Angrist, J. D. and Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.

Athey, S. and Imbens, G. W. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360.

Athey, S., Tibshirani, J., and Wager, S. (2019). Generalized random forests. *The Annals of Statistics*, 47(2):1148 – 1178.

Boyle, J., Zeitz, K., Hoffman, R., Khanna, S., and Beltrame, J. (2013). Probability of severe adverse events as a function of hospital occupancy. *IEEE journal of biomedical and health informatics*, 18(1):15–20.

Breiman, L. (2001). Random forests. *Machine learning*, 45:5–32.

Castagna, F., Xue, X., Saeed, O., Kataria, R., Puius, Y. A., Patel, S. R., Garcia, M. J., Racine, A. D., Sims, D. B., and Jorde, U. P. (2022). Hospital bed occupancy rate is an independent risk factor for covid-19 inpatient mortality: a pandemic epicentre cohort study. *BMJ open*, 12(2):e058171.

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68.

Elsayed, S., Cosker, T., and Grant, A. (2005). Pressure for beds–does it put our orthopaedic patients at risk? *Injury*, 36(1):86–87.

Handel, D. A., Su, Z., Hendry, N., and Mauldin, P. (2018). Inpatient placement: Associations with mortality, cost, and length of stay. *The American journal of managed care*, 24(7):e230–e233.

Hirano, K. and Porter, J. R. (2009). Asymptotics for statistical treatment rules. *Econometrica*, 77(5):1683–1701.

Kitagawa, T. and Tetenov, A. (2018). Who should be treated? empirical welfare maximization methods for treatment choice. *Econometrica*, 86(2):591–616.

Kuntz, L., Mennicken, R., and Scholtes, S. (2015). Stress on the ward: Evidence of safety tipping points in hospitals. *Management Science*, 61(4):754–771.

Lepage, B., Robert, R., Lebeau, M., Aubeneau, C., Silvain, C., and Migeot, V. (2009). Use of a risk analysis method to improve care management for outlying inpatients in a university hospital. *BMJ Quality & Safety*, 18(6):441–445.

Lloyd, J., Elsayed, S., Majeed, A., Kadambande, S., Lewis, D., Mothukuri, R., and Kulkarni, R. (2005). The practice of out-lying patients is dangerous:: A multicentre comparison study of nursing care provided for trauma patients. *Injury*, 36(6):710–713.

Madsen, F., Ladelund, S., and Linneberg, A. (2014). High levels of bed occupancy associated with increased inpatient and thirty-day hospital mortality in denmark. *Health affairs*, 33(7):1236–1244.

Manski, C. F. (2004). Statistical treatment rules for heterogeneous populations. *Econometrica*, 72(4):1221–1246.

Robins, J. M., Rotnitzky, A., and Zhao, L. P. (1994). Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association*, 89(427):846–866.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688.

Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical association*, 46(253):55–67.

Schilling, P. L., Campbell Jr, D. A., Englesbe, M. J., and Davis, M. M. (2010). A comparison of in-hospital mortality risk conferred by high hospital occupancy, differences in nurse staffing levels, weekend admission, and seasonal influenza. *Medical care*, 48(3):224–232.

Sharma, N., Moffa, G., Schwendimann, R., Endrich, O., Ausserhofer, D., and Simon, M. (2022). The effect of time-varying capacity utilization on 14-day in-hospital mortality: a retrospective longitudinal study in swiss general hospitals. *BMC health services research*, 22(1):1551.

Sprivulis, P. C., Da Silva, J.-A., Jacobs, I. G., Jelinek, G. A., and Frazer, A. R. (2006). The association between hospital overcrowding and mortality among patients admitted via western australian emergency departments. *Medical journal of Australia*, 184(5):208–212.

Stowell, A., Claret, P.-G., Sebbane, M., Bobbia, X., Boyard, C., Genre Grandpierre, R., Moreau, A., and de La Coussaye, J.-E. (2013). Hospital out-lying through lack of beds and its impact on care and patient outcome. *Scandinavian journal of trauma, resuscitation and emergency medicine*, 21:1–7.

Stoye, J. (2009). Minimax regret treatment choice with finite samples. *Journal of Econometrics*, 151(1):70–81.

Wager, S. and Athey, S. (2018a). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242.

Wager, S. and Athey, S. (2018b). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242.

Wise, J. (2015). Cutting medical bed occupancy is linked to reduced patient mortality in uk hospital, study shows. *BMJ*, 351.

A Appendix

Table 13: Descriptive Statistics

Variable	Overall	Int.Med.	Surg.
	Mean (Std Dev)	**Mean (Std Dev)**	**Mean (Std Dev)**
Busyness (Int. Med. NCU)	200.89 (43.38)	183.61 (20.41)	183.36 (20.34)
Busyness (Surg. NCU)	212.25 (27.87)	203.75 (17.96)	205.41 (17.93)
Admission Type (5)	6252 (8.41%)	2198 (15.73%)	1222 (13.99%)
Admission Type (2)	40402 (54.34%)	6942 (49.69%)	5843 (66.91%)
Admission Type (1)	27214 (36.60%)	4817 (34.48%)	1642 (18.80%)
Admission Type (4)	243 (0.33%)	8 (0.06%)	19 (0.22%)
Admission Type (9)	188 (0.25%)	2 (0.01%)	3 (0.03%)
Admission Type (8)	56 (0.08%)	5 (0.04%)	4 (0.05%)
Referring Institution (3)	43716 (58.79%)	11364 (81.33%)	8003 (91.64%)
Referring Institution (2)	11351 (15.27%)	1858 (13.30%)	440 (5.04%)
Referring Institution (1)	5231 (7.04%)	728 (5.21%)	244 (2.79%)
Referring Institution (9)	13804 (18.56%)	5 (0.04%)	37 (0.42%)
Referring Institution (8)	208 (0.28%)	8 (0.06%)	3 (0.03%)
Referring Institution (6)	27 (0.04%)	4 (0.03%)	4 (0.05%)
Referring Institution (5)	9 (0.01%)	4 (0.03%)	2 (0.02%)
Referring Institution (4)	9 (0.01%)	1 (0.01%)	2 (0.05%)
Gender (1)	47858 (64.36%)	9123 (65.29%)	5866 (67.17%)
Gender (2)	26497 (35.64%)	4849 (34.71%)	2867 (32.83%)
Year (2018)	9392 (12.63%)	1693 (12.12%)	1049 (12.01%)
Year (2017)	9038 (12.16%)	1728 (12.37%)	1002 (11.47%)
Year (2019)	10215 (13.74%)	1810 (12.95%)	1192 (13.65%)
Year (2016)	8839 (11.89%)	1653 (11.83%)	988 (11.31%)
Year (2011)	5765 (7.75%)	1295 (9.27%)	878 (10.05%)
Year (2013)	7443 (10.01%)	1271 (9.10%)	841 (9.63%)
Year (2012)	6275 (8.44%)	1267 (9.07%)	830 (9.50%)
Year (2015)	8720 (11.73%)	1706 (12.21%)	986 (11.29%)
Year (2014)	8668 (11.66%)	1549 (11.09%)	967 (11.07%)
Swiss (1)	56850 (76.46%)	12328 (88.23%)	7634 (87.42%)
Swiss (2)	17505 (23.54%)	1644 (11.77%)	1099 (12.58%)
Age	67.99 (14.70)	70.14 (14.04)	66.00 (13.94)

Table 14: Descriptive Statistics by NCU

Variable	Overall	Int.Med.	Surg.
	Count (%)	**Count (%)**	**Count (%)**
ICD-10 (I21)	9292 (12.50%)	2719 (19.46%)	363 (4.16%)
ICD-10 (I25)	8734 (11.75%)	3202 (22.92%)	687 (7.87%)
ICD-10 (I35)	4500 (6.05%)	1172 (8.39%)	689 (7.89%)
ICD-10 (I63)	7784 (10.47%)	1963 (14.05%)	217 (2.48%)
ICD-10 (I34)	1415 (1.90%)	251 (1.80%)	230 (2.63%)
ICD-10 (I74)	1137 (1.53%)	157 (1.12%)	169 (1.94%)
ICD-10 (I65)	774 (1.04%)	136 (0.97%)	204 (2.34%)
ICD-10 (I70)	4850 (6.52%)	1168 (8.36%)	274 (3.14%)
ICD-10 (I67)	1338 (1.80%)	123 (0.88%)	449 (5.14%)
ICD-10 (I20)	2619 (3.52%)	493 (3.53%)	99 (1.13%)
ICD-10 (I62)	948 (1.27%)	23 (0.16%)	380 (4.35%)
ICD-10 (I26)	1207 (1.62%)	254 (1.82%)	20 (0.23%)
ICD-10 (I61)	1470 (1.98%)	126 (0.90%)	273 (3.13%)
ICD-10 (I60)	795 (1.07%)	18 (0.13%)	202 (2.31%)
ICD-10 (I72)	971 (1.31%)	80 (0.57%)	268 (3.07%)
ICD-10 (I71)	2550 (3.43%)	51 (0.37%)	772 (8.84%)
ICD-10 (I31)	346 (0.47%)	106 (0.76%)	26 (0.30%)
ICD-10 (I38)	8 (0.01%)	1 (0.01%)	2 (0.02%)
ICD-10 (I27)	400 (0.54%)	124 (0.89%)	1 (0.01%)
ICD-10 (I77)	296 (0.40%)	73 (0.52%)	36 (0.41%)
ICD-10 (I33)	313 (0.42%)	37 (0.26%)	34 (0.39%)
ICD-10 (I36)	83 (0.11%)	15 (0.11%)	9 (0.10%)
ICD-10 (I73)	104 (0.14%)	35 (0.25%)	6 (0.07%)
ICD-10 (I30)	216 (0.29%)	28 (0.20%)	4 (0.05%)
ICD-10 (I24)	99 (0.13%)	34 (0.24%)	4 (0.05%)
ICD-10 (I28)	27 (0.04%)	8 (0.06%)	2 (0.02%)
ICD-10 (I23)	9 (0.01%)	2 (0.01%)	2 (0.02%)
ICD-10 (I37)	39 (0.05%)	3 (0.02%)	4 (0.05%)
ICD-10 (I66)	27 (0.04%)	8 (0.06%)	3 (0.03%)
ICD-10 (I78)	25 (0.03%)	5 (0.04%)	1 (0.01%)
ICD-10 (I64)	40 (0.05%)	12 (0.09%)	1 (0.01%)
ICD-10 (I22)	22 (0.03%)	4 (0.03%)	1 (0.01%)
ICD-10 (C16)	804 (1.08%)	58 (0.42%)	111 (1.27%)
ICD-10 (C34)	4266 (5.74%)	525 (3.76%)	382 (4.37%)
ICD-10 (C71)	1477 (1.99%)	122 (0.87%)	308 (3.53%)
ICD-10 (C49)	650 (0.87%)	58 (0.42%)	123 (1.41%)
ICD-10 (C19)	152 (0.20%)	5 (0.04%)	23 (0.26%)
ICD-10 (C20)	772 (1.04%)	17 (0.12%)	113 (1.29%)
ICD-10 (C44)	496 (0.67%)	25 (0.18%)	130 (1.49%)
ICD-10 (C79)	2034 (2.74%)	145 (1.04%)	319 (3.65%)
ICD-10 (C22)	1970 (2.65%)	35 (0.25%)	452 (5.18%)
ICD-10 (C77)	501 (0.67%)	23 (0.16%)	101 (1.16%)
ICD-10 (C24)	237 (0.32%)	8 (0.06%)	63 (0.72%)
ICD-10 (C78)	2355 (3.17%)	76 (0.54%)	401 (4.59%)
ICD-10 (C41)	287 (0.39%)	25 (0.18%)	33 (0.38%)
ICD-10 (C18)	1180 (1.59%)	32 (0.23%)	122 (1.40%)
ICD-10 (C73)	702 (0.94%)	8 (0.06%)	92 (1.05%)
ICD-10 (C33)	18 (0.02%)	2 (0.01%)	2 (0.02%)
ICD-10 (C45)	280 (0.38%)	38 (0.27%)	32 (0.37%)
ICD-10 (C48)	148 (0.20%)	12 (0.09%)	19 (0.22%)
ICD-10 (C25)	1009 (1.36%)	62 (0.44%)	176 (2.02%)
ICD-10 (C40)	336 (0.45%)	38 (0.27%)	16 (0.18%)
ICD-10 (C15)	575 (0.77%)	29 (0.21%)	81 (0.93%)
ICD-10 (C17)	146 (0.20%)	6 (0.04%)	28 (0.32%)